Theory, Algorithms, and Applications
of
Level Set Methods for Propagating Interfaces [*]

J.A. Sethian
Dept. of Mathematics
University of California, Berkeley

August 15, 1995

In Press: Acta Numerica

Abstract

We review recent work on level set methods for following the evolution of complex interfaces. These techniques are based on solving an initial value partial differential equations for a level set functions, using techniques borrowed from hyperbolic conservation laws. Topological changes, corner and cusp development, and accurate determination of geometric properties such as curvature and normal direction are naturally obtained in this setting. The methodology results in robust, accurate, and efficient numerical algorithms for propagating interfaces in highly complex settings. We review the basic theory and approximations, describe a hierarchy of fast methods, including an extremely fast marching level set scheme for monotonically advancing fronts, based on a stationary formulation of the problem, and discuss extensions to multiple interfaces and triple points. Finally, we demonstrate the technique applied to a series of examples from geometry, material science and computer vision, including mean curvature flow, minimal surfaces, grid generation, fluid mechanics, combustion, image processing, computer vision, and etching, deposition and lithography in the microfabrication of electronic components.

Contents

V Applications 45

1 Introduction

Propagating interfaces occur in a wide variety of settings. As physical entities, they include ocean waves, burning flames, and material boundaries. Less obvious boundaries are equally important, and include shapes against backgrounds, hand-written characters, and iso-intensity contours in images.

The goal of this paper is to describe some recent work on level set methods which attempts to unify these problems and provide a general framework for modeling the evolution of boundaries. Our aim is to review a collection of state-of-the-art details of computational techniques for tracking moving interfaces, and to give some sense of the flavor and breadth of applications.

Level set methods are numerical techniques which offer remarkably powerful tools for understanding, analyzing, and computing interface motion in a host of settings. At their core, they rely on a fundamental shift in how one views moving boundaries; rethinking the Lagrangian geometric perspective and exchanging it for an Eulerian, initial value partial differential equation perspective. Five clear advantages result from this new view of propagating interfaces:

- First, from a theoretical/mathematical point of view, the real complexities of front motion are illuminated, in particular, the role of singularities, weak solutions, shock formation, entropy conditions and topological change in the evolving interface.

- Second, from a numerical perspective, natural and accurate ways of computing delicate quantities emerge, including the ability to build high order advection schemes, compute local curvature in two and three dimensions, track sharp corners and cusps, and handle subtle topological changes of merger and breakage.

- Third, from an implementation point of view, since the approach is based on an initial value partial differential equation, robust schemes result from numerical parameters set at the beginning of the computation. The error is thus controlled by

  1. The order of the numerical method,
  2. The grid spacing Δh,
  3. The time step Δt.

- Fourth, computational adaptivity, both in meshing and in computational labor, is possible, as is a clear path to parallelism.

- Fifth, in the case of monotonically advancing fronts under certain speed laws, in this paper we introduce exceptionally fast methods based on merging narrow band techniques and sorting algorithms.

In this paper, we survey an illustrative subset of past and current applications of level set methods. By no means is this an exhaustive review. A large body of work has been reluctantly skipped in the effort to keep this paper of reasonable length. The interested reader is referred to the many references references will be given throughout the text.

Part I
Level Set Formulations

2 Theory of Front Evolution

Consider a boundary, either a curve in two dimensions or a surface in three dimensions, separating one region from another. Imagine that this curve/surface moves in its normal direction with a known speed function F. Our goal is to track the motion of this interface as it evolves. We are only concerned with the motion of the interface in its normal direction; throughout, we shall ignore tangential motions of the interface.
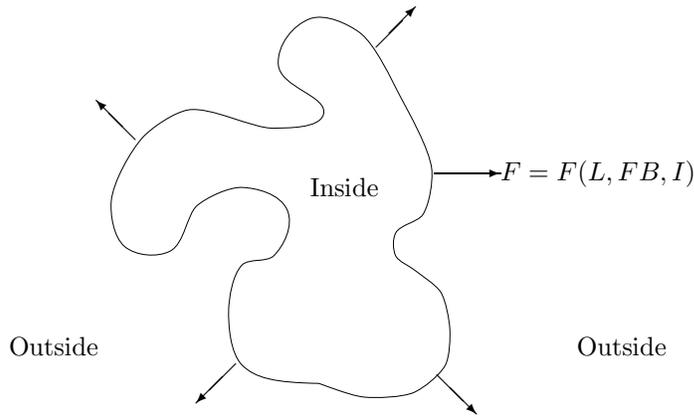
Figure 1: Curve Propagating With Speed F in Normal Direction

The speed function F can be thought of as depending on three types of arguments, namely

$$F = F(L, G, I), \tag{1}$$

where

- $L=$ *Local Properties of the front*, are those determined by local geometric information, such as curvature and normal direction,

- $G=$ *Global Properties of the front*, such as integrals along the front and associated differential equations, are those whose solution depend on the shape and position of the front. For example, suppose the interface is a source of heat which affects diffusion on either side of the interface which in turn influences the motion of the interface. This would be characterized as global argument.

- *I: Independent Properties*, are those which are independent of the shape of the front, such as an underlying fluid velocity which passively transports the front.

Much of the challenge in interface problems comes from producing an adequate model for the speed function F; this is a separate issue independent of the goal of an accurate scheme for advancing the interface based on the model for F. In this section, we assume that the speed function F is known. In Part IV, we discuss the development of models for a collection of applications.

Our first goal is to develop the necessary theory to understand the interplay between the speed function F and the shape of the interface. For ease of discussion, we now turn to the simplest case of a closed curve propagating in the plane.

2.1 Fundamental Formulation

Let γ be a smooth, closed initial curve in R^2, and let $\gamma(t)$ be the one–parameter family of curves generated by moving $\gamma = \gamma(t = 0)$ along its normal vector field with speed F. Here, F is the given scalar function. Thus, we have that $\vec{n} \cdot \vec{x}_t = F$, where $\vec{x}$ is the position vector of the curve, t is time, and $\vec{n}$ is the unit normal to the curve.

As a first attempt, a natural approach is to consider a parameterized description of the motion. We further restrict ourselves and imagine that the speed function F depends only on the local curvature κ of the curve, that is, $F = F(\kappa)$. Thus, we let the position vector $\vec{x}(s,t)$ parameterize γ at time t. Here, $0 \leq s \leq S$, and we assume periodic boundary conditions $\vec{x}(0,t) = \vec{x}(S,t)$. The curve is parameterized so that the interior is on the left in the direction of increasing s (see Fig.2). Let $\vec{n}(s,t)$ be the parameterization of the outward normal and $\kappa(s,t)$ be the parameterization of the curvature. The equations of motion can then be written in terms of individual components $\vec{x} = (x,y)$ as

$$x_t = F[\frac{y_{ss}x_s - x_{ss}y_s}{(x_s^2 + y_s^2)^{3/2}}]~(\frac{y_s}{(x_s^2 + y_s^2)^{1/2}}) \qquad y_t = -F[\frac{y_{ss}x_s - x_{ss}y_s}{(x_s^2 + y_s^2)^{3/2}}]~(\frac{x_s}{(x_s^2 + y_s^2)^{1/2}}),$$

$$(2)$$

where we have used the parameterized expression for the curvature $\kappa = \frac{y_{ss}x_s - x_{ss}y_s}{(x_s^2 + y_s^2)^{3/2}}$ inside the speed function $F(\kappa)$. This is a "Lagrangian" representation because the range of $(x(s,t), y(s,t))$ describes the moving front.
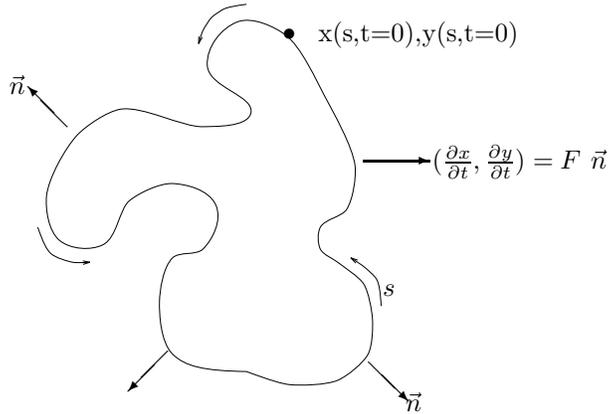
Figure 2: Parameterized View of Propagating Curve

2.2 Total Variation: Stability and the Growth of Oscillations

What happens to oscillations in the initial curve as it moves? We summarize the argument in [106] showing that the decay of oscillations depends only on the sign of F_κ at $\kappa = 0$. The metric $g(s,t)$, which measures the "stretch" of the parameterization, is given by $g(s,t) = (x_s^2 + y_s^2)^{1/2}$. Define the total oscillation (also known as the total variation) in the front

$$Var(t) = \int_0^S |\kappa(s,t)|g(s,t)ds. \qquad (3)$$

This measures the amount of "wrinkling". Our goal is to find out if this wrinkling increases or decreases as the front evolves (see Figure 3).

Differentiation of both the curvature and the metric with respect to time, together with substitution from Eqn. 2 produces the corresponding evolution equations for the metric and curvature, namely

$$\kappa_t = -g^{-1}(F_s g^{-1})_s - \kappa^2 F \tag{4}$$

$$g_t = g\kappa F \tag{5}$$

(Here, g^{-1} is $1/g$, not the inverse). Now, suppose we have a non-convex initial curve moving with speed $F(\kappa)$, and suppose the moving curve stays smooth. By evaluating the time change of the total variation in the solution, we have the following (see [106]):

Original Curve   *Decrease in Variation*   *Increase in Variation*
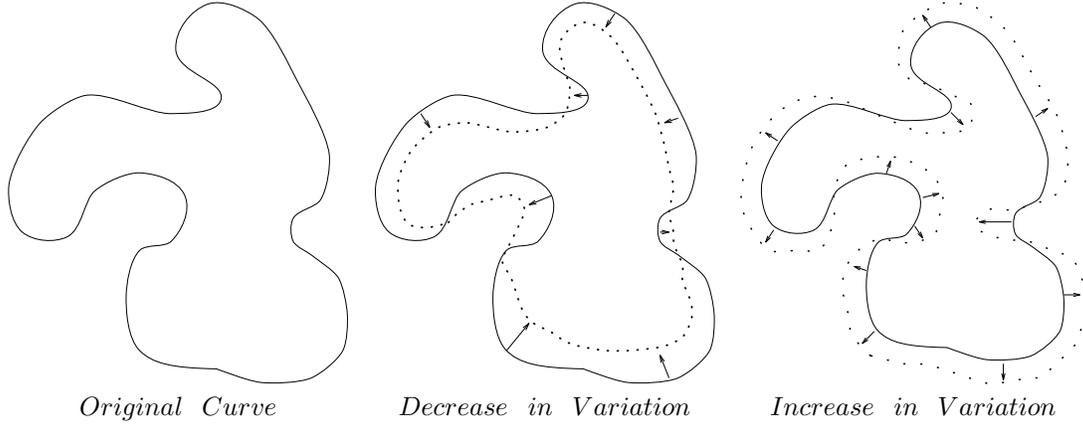
Figure 3: Change in Variation

THEOREM

Consider a front moving along its normal vector field with speed $F(\kappa)$, as in Eqn. 2. Assume that the initial curve $\gamma(0)$ is smooth and non-convex, so that $\kappa(s, 0)$ changes sign. Assume that F is twice differentiable, and that $\kappa(s, t)$ is twice differentiable for $0 \leq s \leq S$ and $0 \leq t \leq T$. Then, for $0 \leq t \leq T$,

- if $F_\kappa \leq 0$ ($F_\kappa \geq 0$) wherever $\kappa = 0$, then

$$\frac{dVar(t)}{dt} \leq 0 \ (\frac{dVar(t)}{dt} \geq 0); \tag{6}$$

- if $F_\kappa < 0$ ($F_\kappa > 0$) and $\kappa_s \neq 0$ wherever $\kappa = 0$, then

$$\frac{dVar(t)}{dt} < 0 \ (\frac{dVar(t)}{dt} > 0). \tag{7}$$

Remarks:

The theorem states that if $F_\kappa < 0$ wherever $\kappa = 0$, then the total variation decreases as the front moves and the front "smooths out", that is, the energy of the front dissipates. The front is assumed to remain smooth in the interval $0 \leq t \leq T$ (the curvature is assumed to be twice differentiable). (In the next section, we discuss what happens if the front ceases to be smooth and develops a corner.) In the special case that $\gamma(t)$ is convex for all t, Proposition 1 is trivial, since $Var(t) = \int_0^S \kappa g ds = 2\pi$. The proof may be found in [106].

7

Swallowtail$(F = 1.0)$ *Entropy Solution*$(F = 1.0)$

Fig.4a   *Fig.4b*
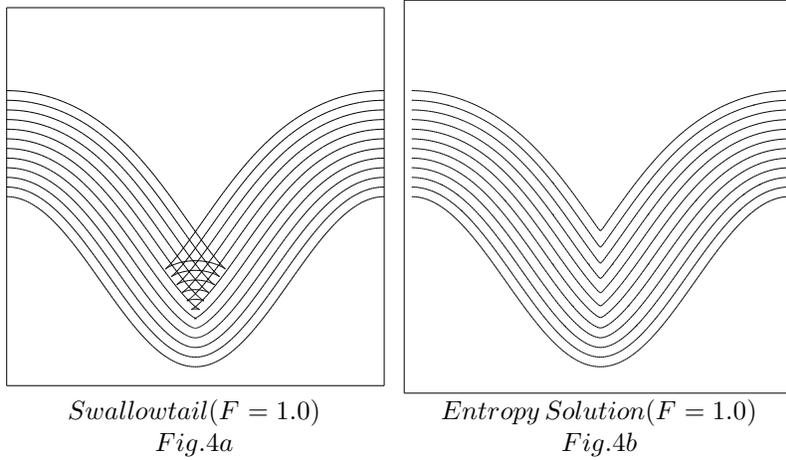
Figure 4: Cosine Curve Propagating with Unit Speed

Two important cases can be easily checked. A speed function $F(\kappa) = 1 - \epsilon\kappa$ for ϵ positive has derivative $F_\kappa = -\epsilon$, and hence the total variation decays. Conversely, a speed function of the form $F(\kappa) = 1 + \epsilon\kappa$ yields a positive speed derivative, and hence oscillations grow. We shall see that the sign of the curvature term in this case corresponds to the backwards heat equation, and hence must be unstable.

2.3 The Role of Entropy Conditions and Weak Solutions

The above theorem assumes that the front stays smooth and differentiable. In many cases of evolving fronts, differentiability is soon lost. For example, consider the periodic initial cosine curve

$$\gamma(0) = (-s, [1 + \cos 2\pi s]/2) \tag{8}$$

propagating with speed $F(\kappa) = 1$. The exact solution to this problem at time t may be easily constructed by advancing each point of the front in its normal direction a distance t. In fact, in terms of our parameterization of the front, the solution is given by

$$x(s, t) = \frac{y_s(s, t = 0)}{(x_s^2(s, t = 0) + y_s^2(s, t = 0))^{1/2}} \, t + x(s, t = 0), \tag{9}$$

$$y(s, t) = -\frac{x_s(s, t = 0)}{(x_s^2(s, t = 0) + y_s^2(s, t = 0))^{1/2}} \, t + y(s, t = 0). \tag{10}$$

In Figure 4, the solution is given for this propagating cosine curve.

It is clear that the front develops a sharp corner, known as a *shock*, in finite time; however, once this corner develops, it is not at all clear how to construct the normal at the corner and continue the evolution, since the derivative is not defined there. Thus, beyond the formation of the discontinuity in the derivative, we need a *weak solution*, so-called because the solution weakly satisfies the definition of differentiability

How shall we continue the solution beyond the formation of the singularity in the curvature corresponding to the corner in the front? The correct answer depends on the nature of the interface under discussion. If we regard the interface as a geometric curve evolving under the prescribed speed function, then one possible weak solution is the "swallowtail" solution formed by letting the front pass through itself; this is the solution shown in Figure 4a. We note that this solution is in fact the one given by Eqns. 9, 10;

Shock   Rarefaction Fan
Fig.5a   Fig.5b
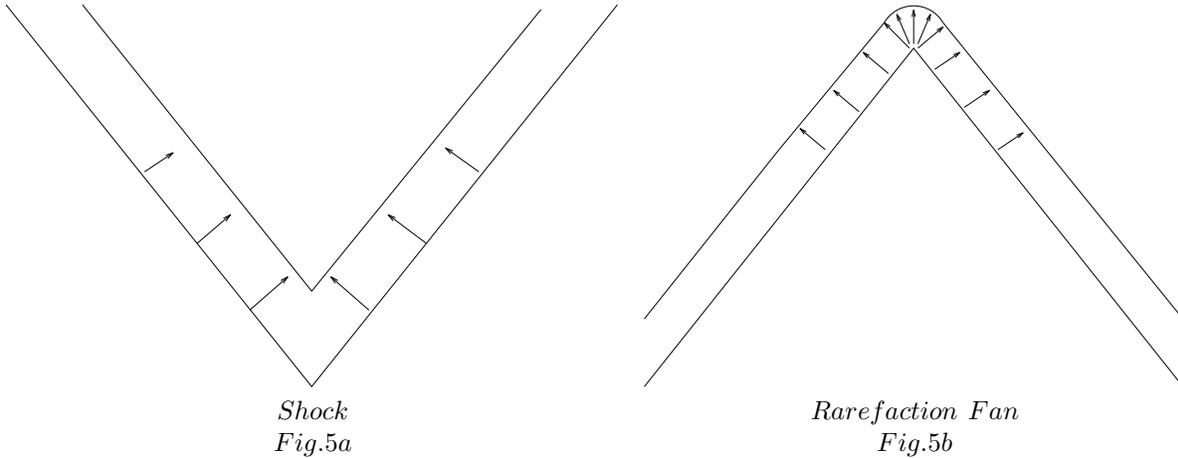
Figure 5: Front Propagating with Unit Normal Speed

the lack of differentiability at the center point does not destroy the solution, since we have written the solution in terms of the initial data.

However, if we regard the moving curve as an interface separating two regions, the front at time t should consist of only the set of all points located a distance t from the initial curve. (This is known as the Huygens principle construction, see [106]). Roughly speaking, we want to remove the "tail" from the "swallowtail". In Figure 4b, we show this alternate weak solution. Another way to characterize this weak solution is through the following "*entropy condition*" posed by Sethian (see [106, 105]): If the front is viewed as a burning flame, then *once a particle is burnt it stays burnt.* Careful adherence to this stipulation produces the Huygens principle construction.

What does this "entropy condition" have to do with the notion of "entropy"? An intuitive answer is that an entropy condition stipulates that no new information can be created during the evolution of the problem. Once an entropy condition is invoked, some information about the initial data is lost. Indeed, our entropy condition says that once a particle is burnt, it stays burnt, that is, once a corner has developed, the solution is no longer reversible. The problem cannot be run "backwards" in time; if we try to do so, we will not retrieve the initial data. Thus, some information about the solution is forever lost.

As further illustration, we consider the case of a V-shaped front propagating normal to itself with unit speed ($F = 1$). In Figure 5a, the point of the front is downwards; as the front moves inwards with unit speed, a shock develops as the front pinches off, and an entropy condition is required to select the correct solution to stop the solution from being multiple-valued. Conversely, in Figure 5b, the point of the front is upwards; in this case the unit normal speed results in a circular fan which connects the left state with slope $+1$ to the right state which has slope -1.

It is important to summarize a key point in the above discussion. Our choice of weak solution given by our entropy condition rests on the perspective that the front separates two regions, and the assumption that we are interested in tracking the progress of one region into the other.

$$F = 1. - 0.25\kappa \qquad\qquad F = 1. - \epsilon\kappa \qquad\qquad Entropy\,Solution(F = 1.0)$$

$$Fig.6a. \qquad\qquad\qquad\qquad\qquad\qquad\qquad\qquad Fig.6b$$
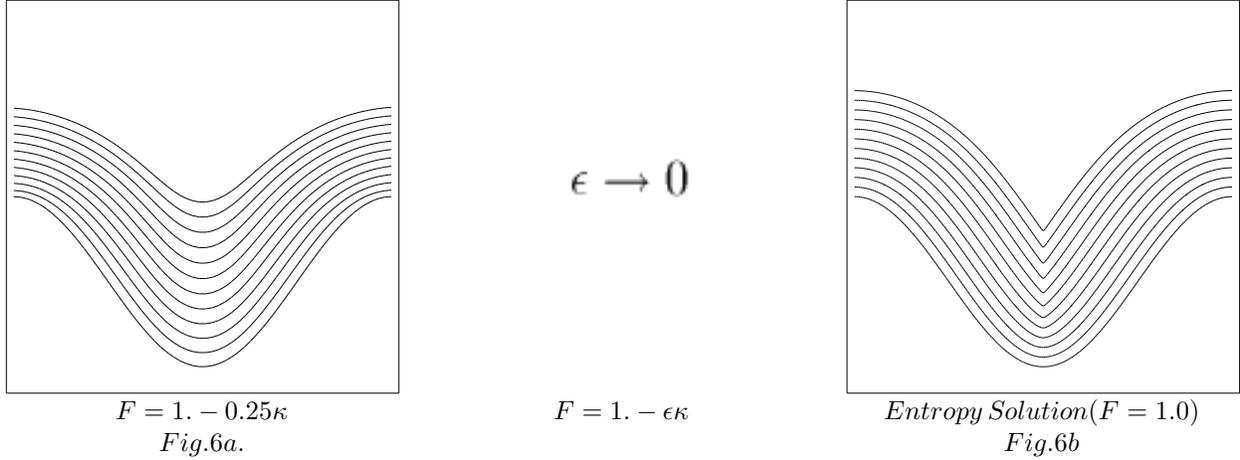
Figure 6: Entropy-Solution is the Limit of Viscous Solution

2.4   Effects of Curvature: The Viscous Limit and the Link to Hyperbolic Conservation Laws

Consider now a speed function of the form $F = 1 - \epsilon\kappa$, where ϵ is a constant. The modifying effects of curvature on the former constant speed law are profound, and in fact pave the way towards constructing accurate numerical schemes which adhere to the correct entropy condition.

Following [106], the curvature evolution equation given by Eqn. 4 can be rewritten in terms of arclength, namely

$$\kappa_t = \epsilon\kappa_{\alpha\alpha} + \epsilon\kappa^3 - \kappa^2, \tag{11}$$

where the second derivative of the curvature κ is taken with respect to arclength α. This is a reaction-diffusion equation; the drive toward singularities due to the reaction term $(\epsilon\kappa^3 - \kappa^2)$ is balanced by the smoothing effect of the diffusion term $(\epsilon\kappa_{\alpha\alpha})$. Indeed, with $\epsilon = 0$, we have a pure reaction equation $\kappa_t = -\kappa^2$, and the developing corner can be seen in the exact solution $\kappa(s,t) = \kappa(s,0)/(1 + t\kappa(s,0))$, which is singular in finite t if the initial curvature is anywhere negative. Thus, as shown, corners can form in the moving curve when $\epsilon = 0$.

Consider again the cosine front given in Eqn. 8 and the speed function $F(\kappa) = 1 - \epsilon\kappa$, $\epsilon > 0$. As the front moves, the troughs at $s = n + 1/2, n = 0, \pm 1, \pm 2, ...$ are sharpened by the negative reaction term (because $\kappa < 0$ at such points) and smoothed by the positive diffusion term (see Figure 6a). For $\epsilon > 0$, it can be shown (see [106, 88]) that the moving front stays C^∞. The entropy-satisfying solution to this problem when $F = 1$ from Figure 4b is shown in Figure 6b.

The central observation, given in [106] and key to the level set approach, is the following link:

Consider the above propagating cosine curve and the two solutions:

- $X^\epsilon_{curvature}(t)$, *obtained by evolving the initial front with* $F_\epsilon = 1 - \epsilon\kappa$,

- $X_{constant}(t)$, *obtained with speed function* $F = 1$ *and the entropy condition,*

Then, at any time T,

$$\lim_{\epsilon \to 0} X^\epsilon_{curvature}(T) = X_{constant}(T). \tag{12}$$
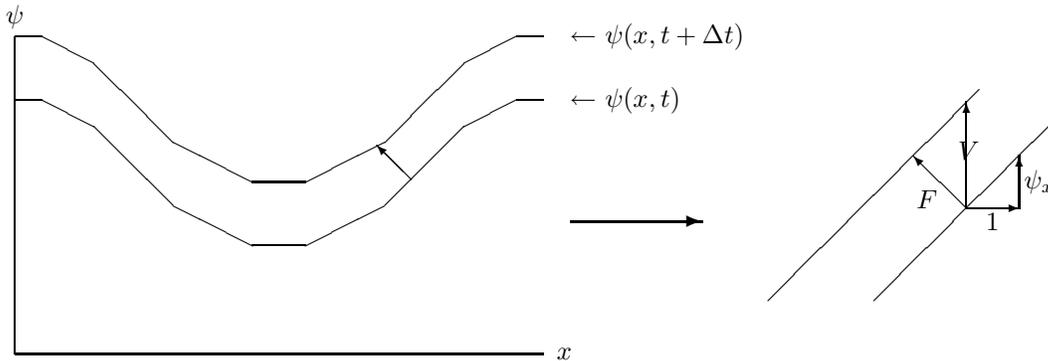
Figure 7: Variables For Propagating Graph

Thus, the limit of motion with curvature, known as the "viscous limit", is the entropy solution for the constant speed case, see [106].

Why is this known as the "viscous limit", or, more accurately, what does this have to do with viscosity? In order to see why viscosity is an appropriate name, we turn to the link between propagating fronts and hyperbolic conservation laws; for details, see [106].

An equation of the form

$$u_t + [G(u)]_x = 0 \qquad (13)$$

is known as a hyperbolic conservation law. A simple example is Burger's equation, given by

$$u_t + uu_x = 0 \qquad (14)$$

which describes the motion of a compressible fluid in one-dimension. The solution to this equation can develop discontinuities, known as "shocks", where the fluid undergoes a sudden expansion or compression. These shocks (for example, a sonic boom) can arise from arbitrarily smooth initial data; they are a function of the equation itself. However, if one includes the effects of fluid viscosity, the equation includes a right-hand-side, namely

$$u_t + uu_x = \epsilon u_{xx}. \qquad (15)$$

This second derivative on the right-hand-side acts like a smoothing term and stops the development of such shocks; it can be shown that the solutions must remain smooth for all time.

What does this have to do with our propagating front equation? Consider the initial front given by the graph of $f(x)$, with f periodic on $[0, 1]$, and suppose that the propagating front remains a function for all time. Let ψ be the height of the propagating function at time t, thus $\psi(x, 0) = f(x)$. The tangent at (x, ψ) is $(1, \psi_x)$. Referring to Figure 7, the change in height V in a unit time is related to the speed F in the normal direction by

$$\frac{V}{F} = \frac{(1 + \psi_x^2)^{1/2}}{1}, \qquad (16)$$

11

and thus the equation of motion becomes

$$\psi_t = F(1 + \psi_x^2)^{1/2}. \tag{17}$$

Using the speed function $F(\kappa) = 1 - \epsilon\kappa$ and the formula $\kappa = -\psi_{xx}/(1 + \psi_x^2)^{3/2}$, we get

$$\psi_t \ - \ (1 + \psi_x^2)^{1/2} = \epsilon\frac{\psi_{xx}}{1 + \psi_x^2}. \tag{18}$$

We first note that this is a partial differential equation with a first order time and space derivative on the left hand side, and a parabolic second-order term on the right. Differentiating both sides of this equation yields an evolution equation for the slope $u = d\psi/dx$ of the propagating front, namely

$$u_t + [-(1 + u^2)^{1/2}]_x = \epsilon[\frac{u_x}{1 + u^2}]_x. \tag{19}$$

Thus, as shown in [113], the derivative of our equation with parabolic right-hand-side for the changing height ψ looks like a viscous hyperbolic conservation law with $G(u) = (1 + u^2)^{1/2}$ for the propagating slope u (see [113]). Hyperbolic conservation laws of the above form have a long history, in fact, our entropy condition is equivalent to the one for propagating shocks in hyperbolic conservation laws (see [106]). Our goal will be to exploit the theory and technology of numerical solutions of hyperbolic conservation laws to devise accurate numerical schemes to solve the equation of motion for propagating fronts.

Before doing so, however, we have a technical problem. The equation of motion given by Eqn. 17 only refers to fronts which remain the graph of a function as they move. The above ideas must be extended to include propagating fronts which are not easily written as functions. This is the time-dependent level set idea introduced by Osher and Sethian [88].

3 The Time-Dependent Level Set Formulation

3.1 Formulation

Given a closed $N - 1$ dimensional hypersurface $\Gamma(t)$, we now produce an *Eulerian* formulation for the motion of the hypersurface propagating along its normal direction with speed F, where F can be a function of various arguments, including the curvature, normal direction, etc. The main idea of the level set methodology is to embed this propagating interface as the zero level set of a higher dimensional function ϕ. Let $\phi(x, t = 0)$, where x is a point in R^N, be defined by

$$\phi(x, t = 0) = \pm d, \tag{20}$$

where d is the distance from x to $\Gamma(t = 0)$, and the plus (minus) sign is chosen if the point x is outside (inside) the initial hypersurface $\Gamma(t = 0)$. Thus, we have an initial function $\phi(x, t = 0) : R^N \to R$ with the property that

$$\Gamma(t = 0) = (x|\phi(x, t = 0) = 0). \tag{21}$$

Our goal is to produce an equation for the evolving function $\phi(x, t)$ which contains the embedded motion of $\Gamma(t)$ as the level set $\phi = 0$. Let $x(t)$ be the path of a point on the propagating front. That is, $x(t = 0)$ is a point on the initial front $\Gamma(t = 0)$, and $x_t \cdot n = F(x(t))$ where n is the normal to the front at $x(t)$. Since we want the zero level

12
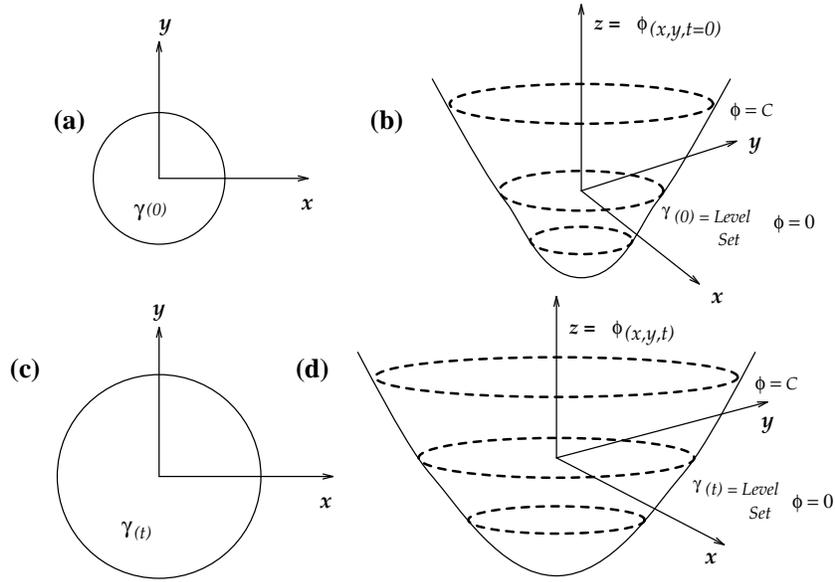
Figure 8: Propagating Circle

set of the evolving function ϕ to always match the propagating hypersurface, we must have

$$\phi(x(t), t) = 0. \tag{22}$$

By the chain rule,

$$\phi_t + \nabla\phi(x(t), t) \cdot x'(t) = 0. \tag{23}$$

Since $n = \nabla\phi/|\nabla\phi|$, we have the evolution equation for ϕ, namely

$$\phi_t + F|\nabla\phi| = 0 \tag{24}$$

$$\phi(x, t = 0) \quad given. \tag{25}$$

This is our time-dependent level set equation. For certain forms of the speed function F, we obtain a standard Hamilton-Jacobi equation.

In Figure 8 (taken from [108]), we show the outward propagation of an initial curve and the accompanying motion of the level set function ϕ.

In Figure 8a, we show the initial circle, and in Figure 8c, we show the circle at a later time. In Figure 8b, we show the initial position of the level set function ϕ, and in Figure 8d, we show this function at a later time. We refer to this as an Eulerian formulation because the underlying coordinate system remains fixed.

3.2 Advantages

There are four major advantages to this Eulerian level set formulation.

1. First, the evolving function $\phi(x, t)$ always remains a function as long as F is smooth. However, the level surface $\phi = 0$, and hence the propagating hypersurface $\Gamma(t)$, may change topology, break, merge, and form sharp corners as the function ϕ evolves (see [88]).

2. The second major advantage of this Eulerian formulation concerns numerical approximation. Because $\phi(x,t)$ remains a function as it evolves, we may use a discrete grid in the domain of x and substitute finite difference approximations for the spatial and temporal derivatives. For example, using a uniform mesh of spacing h, with grid nodes (i,j), and employing the standard notation that ϕ_{ij}^n is the approximation to the solution $\phi(ih, jh, n\Delta t)$, where Δt is the time step, we might write

$$\frac{\phi_{ij}^{n+1} - \phi_{ij}^n}{\Delta t} + (F)(\nabla_{ij}\phi_{ij}^n) = 0 \tag{26}$$

Here, we have used forward differences in time, and let $\nabla_{ij}\phi_{ij}^n$ be some appropriate finite difference operator for the spatial derivative. Thus, an explicit finite difference approach is possible.

3. The third major advantage of the above formulation is that intrinsic geometric properties of the front may be easily determined from the level set function ϕ. For example, at any point of the front, the normal vector is given by

$$\vec{n} = \frac{\nabla\phi}{|\nabla\phi|}, \tag{27}$$

and the curvature of each level set is easily obtained from the divergence of the gradient of the unit normal vector to the front, i.e.,

$$\kappa = \nabla \cdot \frac{\nabla\phi}{|\nabla\phi|} = \frac{\phi_{xx}\phi_y^2 - 2\phi_x\phi_y\phi_{xy} + \phi_{yy}\phi_x^2}{(\phi_x^2 + \phi_y^2)^{3/2}}. \tag{28}$$

4. Finally, the fourth major advantage of the above level set approach is that there are no significant changes required to follow fronts in three dimensions. By simply extending the array structures and gradient operators, propagating surfaces are easily handled.

As an example of the application of level set methods, consider once again the problem of a front propagating with speed $F(\kappa) = 1 - \epsilon\kappa$. In Figure 9, we show two cases of a propagating initial triple sine curve. For ϵ small (Fig. 9a), the troughs sharpen up and result in transverse lines that come too close together. For ϵ large (Fig. 9b), parts of the boundary with high values of positive curvature can initially move downwards, and concave parts of the front can move quickly upwards.

3.3   Theoretical Aspects of the Level Set Formulation

Numerical techniques for approximating moving fronts have been the focus of much effort in computational physics. At the same time, the theoretical analysis of moving curves and surfaces has been a subject of considerable importance in its own right. The work of Gage [48], Gage and Hamilton [49], and Grayson [52], discussed later provided some groundbreaking analysis of flow of curves under the curvature, and the seminal result that a closed curve shrinking under its curvature collapses smoothly to a point.

Along a very different approach, Brakke [15] applied varifold theory to the problem of a hypersurface moving under its curvature, and in doing so provided a wide-ranging

$$F = 1. - 0.025\kappa$$
$$Fig.9a.$$

$$F = 1. - 0.25\kappa$$
$$Fig.9b$$
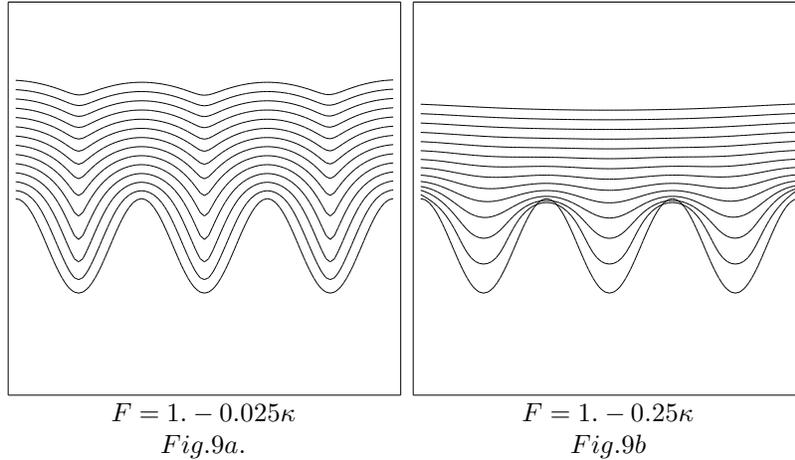
Figure 9: Propagating Triple Sine Curve.

perspective for these problems. His was a general approach, and included problems in which the results were not necessarily smooth. His analysis provided a detailed look at surface curvature evolution problems.

There has been considerable theoretical analysis of the level set approach, its formulation, and its relation to other perspectives on front propagation. The model posed in [105], which considered flame propagation as a function of curvature, and introduced an entropy condition for evolving fronts, served as the basis for theoretical analysis by Barles [8]. The full level set methodology of [88] provides a view of surface evolution different from the one provided in the work of Gage, Grayson, and Brakke. First, the embedding of the front as a higher dimensional function naturally accounts for some of the issues of topological change and corner formation. Second, and more importantly, the transformation of a geometry problem into an initial value partial differential equation means that available technology, including regularity of solutions, viscous solutions of Hamilton-Jacobi equations, and questions of existence and uniqueness, can be applied in this geometrical setting.

Using the above level set approach, Evans and Spruck [41, 42, 43, 44], and Chen, Giga, Goto and Ishii [26, 50, 51] performed detailed analysis of curvature flow in a series of papers. They exploited much of the work on viscosity solutions of partial differential equations developed over the past 15 years (see [71]), which itself was inspired by the corresponding work applied to hyperbolic conservation laws. These papers examined the regularity of curvature flow equations, pathological cases, and the link between the level set perspective and the varifold approach of Brakke. These papers opened up a series of investigations into further issues; we also refer the interested reader to [59, 60, 40].

3.4 Summary

The discussion in Part I may be summarized as follows:

1. A front propagating at a constant speed can form corners as it evolves; at such points, the front is no longer differentiable and a weak solution must be constructed to continue the solution.

2. The correct weak solution, motivated by viewing the front as an evolving interface separating two regions, comes by means of an entropy condition.

3. A front propagating at a speed which depends on its curvature does not form corners and stays smooth for all time. Furthermore, the limit of this motion as the dependence on curvature vanishes is the entropy-satisfying solution obtained for the constant speed case.

4. If the propagating front remains a graph as it moves, there is a direct link between the equation of motion and one-dimensional hyperbolic conservation laws. The role of curvature in a propagating front is analogous to the role of viscosity in equations of viscous compressible fluid flow.

5. By embedding the motion of a curve as the zero level set of a higher dimensional function, an initial value partial differential equation can be obtained which extends the above to include arbitrary curves and surfaces moving in two and three space dimensions.

4   The Stationary Level Set Formulation

In the above level set equation

$$\phi_t + F|\nabla \phi| = 0 \tag{29}$$

the position of the front is given by the zero level set of ϕ at a time t. Suppose we now restrict ourselves to the particular case of a front propagating with a speed F that is either always positive or always negative. In this case, we can convert our level set formulation from a time-dependent partial differential equation to a stationary one in which time has disappeared. We now describe a stationary level set formulation, which is common in control theory.

To explain this transformation, imagine the two-dimensional case in which the interface is a propagating curve, and suppose we graph the evolving zero level set above the xy plane. That is, let $T(x, y)$ be the time at which the curve crosses the point (x, y). The surface $T(x, y)$ then satisfies the equation

$$|\nabla T|F = 1. \tag{30}$$

Eqn. 30 simply says that the gradient of arrival time surface is inversely proportional to the speed of the front. This is a Hamilton-Jacobi equation, and the recasting of the a front motion problem into a stationary one is common in a variety of applications, see [45, 46]. In the case where the speed function F depends only on position, we get the well-known Eikonal equation. The requirement that the speed function always be positive[1] is so that the crossing time surface $T(x, y)$ is single-valued.

To summarize,

- In the time-dependent level set equation, the position of the front Γ at time t is given by the zero level set of ϕ at time t, that is $\Gamma(t) = \{(x, y)|\phi(x, t) = 0\}$.

---

[1] or conversely, always negative

- In the stationary level set equation, the position of the front Γ is given by the level set of value t of the function $T(x, y)$, that is $\Gamma(t) = \{(x, y) | T(x, y) = t\}$.

That is, we wish to solve

Time-Dependent Formulation   **Stationary Formulation**

$$\phi_t + F|\nabla\phi| = 0 \qquad\qquad |\nabla T|F = 1 \qquad\qquad (31)$$

Front= $\Gamma(t) = \{(x, y) | \phi(x, t) = 0\}$ Front= $\Gamma(t) = \{(x, y) | T(x, y) = t\}$

Applies for arbitrary F $\qquad\qquad$ Requires $F > 0$

In both cases, we require an "entropy-satisfying" approximation to the gradient term. In this next section, we discuss appropriate approximations for this term, leading to schemes for both the time-dependent and stationary level set formulations. Our goal now is to turn to issue of numerical approximations, and to develop the necessary theory and numerics to accurately approximate the level set initial value partial differential equation.

Part II
Numerical Approximation

5 Traditional Techniques for Tracking Interfaces

Before discussing the numerical approximation of these level set equations, it is instructive to review briefly more traditional techniques for computing the motion of interfaces.

- *Marker/String Methods:* In these methods, a discrete parametrized version of the interface boundary is used. In two dimensions, marker particles are used; in three dimensions, a nodal triangularization of the interface is often developed. The positions of the nodes are then updated by determining front information about the normals and curvature from the marker representation. Such representations can be quite accurate, however, limitations exist for complex motions. To begin, if corners and cusps develop in the evolving front, markers usually form "swallowtail" solutions which must be removed through delooping techniques which attempt to enforce an entropy condition inherent in such motion (see [106]). Second, topological changes are difficult to handle; when regions merge, some markers must be removed. Third, significant instabilities in the front can result, since the underlying marker particle motions represent a weakly ill-posed initial value problem (see [88]). Finally, extensions of such methods to three dimensions require additional work.

- *Cell-Based Methods:* In these methods, introduced by Noh and Woodward [85], the computational domain is divided into a set of cells which contain "volume fractions". These volume fractions are numbers between 0 and 1, and represent the fraction of each cell containing the physical material. At any time, the front can be reconstructed from these volume fractions. Since their introduction, many elaborate reconstruction techniques have been developed over the years to include pitched slopes and curved surfaces, see [30, 57, 93]. The accompanying accuracy depends on the sophistication of the reconstruction and the so-called "advection sweeps". Some of the most elaborate and accurate versions of these schemes to date are due to Puckett, see [93]. Advantages of such techniques include the ability to easily handle topological changes, design adaptive mesh methods, and extend the results to three dimensions. However, determination of geometric quantities such as normals and curvature can be inaccurate.

- *Characteristic Methods:* In these methods, "ray-trace"-like techniques are used. The characteristic equations for the propagating interface are used, and the entropy condition at forming corners (see [106]) is formally enforced by constructing the envelope of the evolving characteristics. Such methods handle the looping problems more naturally, but may be complex in three-dimensions and require adaptive adding and removing rays, which can cause instabilities and/or oversmoothing.

6 A First Attempt at Constructing An Approximation to the Gradient

We now turn to our time-dependent level set equation itself, and attempt to construct a numerical method.

Recall that our goal is to solve the equation given in Eqn. 24 by

$$\phi_t + F|\nabla\phi| = 0, \tag{32}$$

$$\phi(x, t = 0) \quad given. \tag{33}$$

The marker particle method discretizes the front. The volume-of-fluid (VOF) method divides the domain space into cells which contain fractions of material. The level set method divides the domain up into grid points which discretize the values of the level set function ϕ. Thus, the grid values give the height of a surface above the domain, and if we slice this surface by the xy plane, we extract the zero level set corresponding to the front.

Another way to look at this is to say that each grid point contains the value of the level set function at that point. Thus, there is an entire family of contours, only one of which is the zero level set (see Figure 10). Rather than move each of the contours in a Lagrangian fashion, we stand at each grid point and update its value to correspond to the motion of the surface, thus producing a new contour value at that grid point.
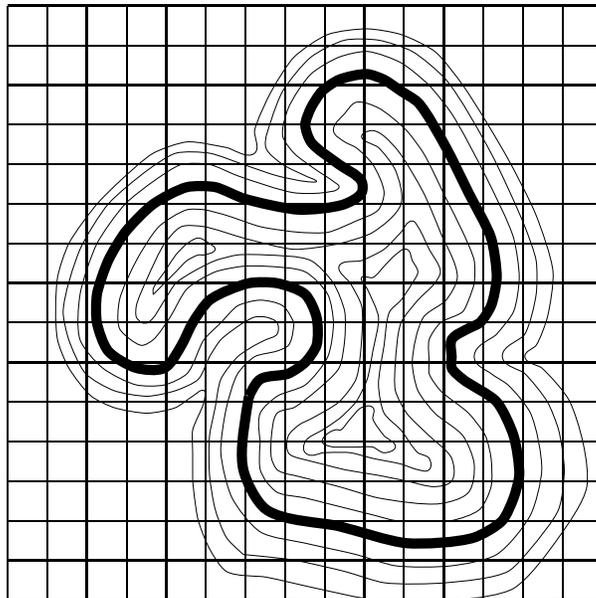
Figure 10: Dark Line is Zero Level Set Corresponding to Front

What is the right way to approximate this equation? We shall investigate the most straightforward numerical approach we can think of by studying the simpler case of an evolving curve whose position can always be described as the graph of a function. The equation for this case was given in Eqn. 17 as shown in Figure 7, namely

$$\psi_t = F(1 + \psi_x^2)^{1/2}. \tag{34}$$

Perhaps the most straightforward way of creating an algorithm to approximate the solution to this equation is to replace all spatial derivatives with central differences and the time derivative with a forward difference, just as we did in in the Lagrangian case.

However, it is easy to see that such an algorithm may not work. Let $F(\kappa) = 1$ and consider the initial value problem

$$\psi_t = (1 + \psi_x^2)^{1/2}, \qquad (35)$$

$$\psi(x, 0) = f(x) = \left\{ \begin{array}{ll} 1/2 - x & x \le 1/2 \\ x - 1/2 & x > 1/2 \end{array} \right\}. \qquad (36)$$

The initial front is a "V" formed by rays meeting at $(1/2, 0)$. By our entropy condition, the solution at any time t is the set of all points located a distance t from the initial "V". To construct a numerical scheme, divide the interval $[0, 1]$ into $2M - 1$ points, and form the central difference approximation to the spatial derivative ψ_x in Eqn. 35, namely

$$\frac{\psi_i^{n+1} - \psi_i^n}{\Delta t} = [1 + [\frac{\psi_{i+1}^n - \psi_{i-1}^n}{2\Delta x}]^2]^{1/2} \qquad (37)$$

Since $x_M = 1/2$, by symmetry, $\psi_{M+1} = \psi_{M-1}$, thus $\psi_t(1/2, 0) = 1$. However, for all $x \ne 1/2$, ψ_t is correctly calculated to be $\sqrt{2}$, since the graph is linear on either side of the corner and thus the central difference approximation is exact. Note that this has nothing to do with the size of the space step Δx or the time step Δt. *No matter how small we take the numerical parameters, as long as we use an odd number of points, the approximation to ψ_t at $x = 1/2$ gets no better.* It is simply due to the way in which the derivative ψ_x is approximated. In Figure 11, we show results using this scheme, with the time derivative ψ_t replaced by a forward difference scheme.

Exact Solution   *Central Differences $\Delta t = .005$*   *Central Differences $\Delta t = .0005$*
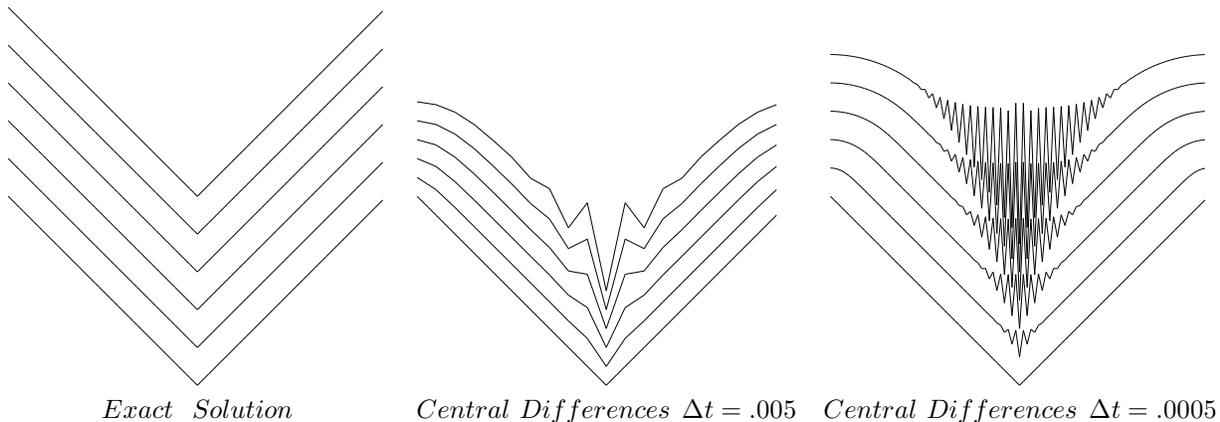
Figure 11: Central Difference Approximation to Level Set Equation

It is easy to see what has gone wrong. In the exact solution, $\psi_t = \sqrt{2}$ for all $x \ne 1/2$. This should also hold at $x = 1/2$ where the slope is not defined; the Huygens construction sets $\psi_t(x = 1/2, t)$ equal to $\lim_{x \to 1/2} \psi_t$. Unfortunately, the central difference approximation chooses a different (and, for our purpose, wrong) limiting solution. It sets the undefined slope ψ_x equal to the average of the left and right slopes. As the calculation progresses, this miscalculation of the slope propagates outwards from the spike as wild oscillations. Eventually, these oscillations cause blowup in the solution.

It is clear that some more care must be taken in formulating an algorithm. What we require are schemes that approximate the gradient term $|\nabla \phi|$ in a way that correctly accounts for the entropy condition. This is the topic of the next section.

7 Schemes from Hyperbolic Conservation Laws

Our schemes are linked to those from hyperbolic conservation laws. As motivation, consider the single scalar hyperbolic conservation law

$$u_t + [G(u)]_x = 0. \tag{38}$$

It is well-known that discontinuities known as shocks can develop in the solution, even with smooth initial data, see [68, 70]. These discontinuities occur because of the collision of characteristics, and an appropriate weak solution must be constructed to carry the solution beyond the collision time. The correct "entropy-solution" is obtained by considering the limit of the associated viscous conservation laws $u_t + [G(u)]_x = \epsilon u_{xx}$ as the viscous coefficient ϵ goes to zero.

From a numerical point of view, the equation can be approximated through the construction of appropriate numerical fluxes g such that

$$\frac{u_i^{n+1} - u_i^n}{\Delta t} = -\frac{g(u_i^n, u_{i+1}^n) - g(u_{i-1}^n, u_i^n)}{\Delta x}. \tag{39}$$

where we require that $g(u, u) = G(u)$. A wide collection of numerical flux functions are available, such as the Lax-Friedrichs flux, Godunox flux, and TVD schemes, see [34, 70]. The goal in the construction of such flux functions is to make sure that the conservation form of the equation is preserved, make sure the entropy condition is satisfied, and try to give smooth (highly accurate) solutions away from the discontinuities. One of the most straightforward approximate numerical fluxes is the Engquist-Osher scheme [39], which is given by

$$g_{EO}(u_1, u_2) = G(u_1) + \int_{u_1}^{u_2} \min(\frac{dG}{du}, 0)du. \tag{40}$$

For the Burger's equation in which $G(u) = u^2$, we have the particularly compact representation of this flux function as

$$g_{EO}(u_1, u_2) = (\max(u_1, 0)^2 + \min(u_2, 0)^2) \tag{41}$$

This flux function will serve as our core technique for approximating the level set equation.

8 Approximating the Time-Dependent Level Set Equation

In this section, we develop schemes for the level set equation

$$\phi_t + F|\nabla\phi| = 0. \tag{42}$$

We begin by writing this equation with a general Hamiltonian H as

$$\phi_t + H(\phi_x, \phi_y, \phi_z) = 0. \tag{43}$$

where

$$H(u, v, w) = \sqrt{u^2 + v^2 + w^2}. \tag{44}$$

We begin with the one-dimensional version, that is, $\phi_t + H(\phi_x) = 0$, where $H(u) = \sqrt{u^2}$.
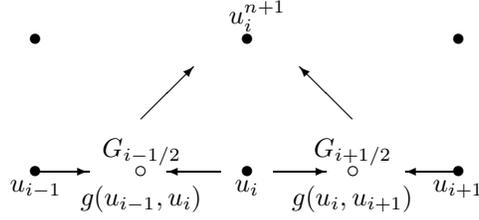
Figure 12: Update of u Through Numerical Flux Function

From the previous section, we have numerical flux functions for the conservation equation $u_t + [G(u)]_x = 0$ which satisfy

$$\frac{u_i^{n+1} - u_i^n}{\Delta t} = -\frac{g(u_i^n, u_{i+1}^n) - g(u_{i-1}^n, u_i^n)}{\Delta x}. \tag{45}$$

In terms of our computational grid shown in Fig.12a, the value of G at the point $(i - 1/2)\Delta x$ (called $G_{i-1/2}$) is approximated by the numerical flux function g as

$$G_{i-1/2} = g(u_{i-1}^n, u_i^n). \tag{46}$$

Similarly, at the point $i + 1/2$, we have

$$G_{i+1/2} = g(u_i^n, u_{i+1}^n). \tag{47}$$

Then from Fig.12, we see that the right-hand-side of Eqn. 39 is just the central difference operator applied to the numerical flux function g. As the grid size goes to zero, consistency requires that $g(u, u) = G(u)$.

We are now all set to build a scheme for our level set equation. Let $u = \phi_x$. Then we can write

$$\phi_t + H(u) = 0. \tag{48}$$

In terms of our computational grid in Fig.13, in order to produce ϕ_i^{n+1} we need ϕ_i^n as well as a value for $H(u_i^n)$. Fortunately, an approximate value for $H(u_i^n)$ is *exactly* what is given by our numerical flux function, hence we have

$$H(u) \approx g(u_{i-1/2}, u_{i+1/2}) \tag{49}$$

All that remains is to construct values for u in the middle of our computational cells. Since $u = \phi_x$, we can use a central difference approximation in ϕ to construct those values. Thus, (see Figure 13), we have

$$\phi_i^{n+1} = \phi_i^n - \Delta t \, g\left(\frac{\phi_i^n - \phi_{i-1}^n}{\Delta x}, \frac{\phi_{i+1}^n - \phi_i^n}{\Delta x}\right) \tag{50}$$

where g is one of the numerical flux functions, and again, we have substituted forward and backward difference operators on ϕ for the values of u at the left and right states.

In the specific case of our one-dimensional level set equation with $H(u) = \sqrt{u^2}$, we can use the EO scheme given in the previous section and, for speed $F = 1$, write

$$\phi_i^{n+1} = \phi_i^n - \Delta t \, (\max(D_i^{-x}, 0)^2 + \min(D_i^{+x}, 0)^2)^{1/2} \tag{51}$$
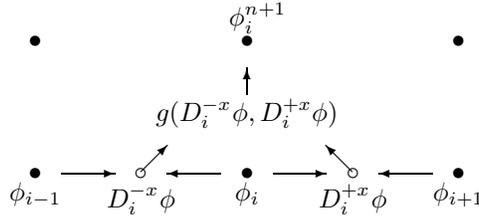
Figure 13: Update of ϕ Through Numerical Hamiltonian

This is the level set scheme given in [88]. As long as the Hamiltonian is symmetric in each of the space dimensions, the above can be replicated in each space variable to construct schemes for two and three-dimensional front propagation problems.

In general then, we follow the following philosophy;

1. If the Hamiltonian "H" is convex, then we use the variant on the Engquist-Osher Scheme presented in [88] and described below.

2. If the Hamiltonian "H" is non-convex, then we use a variant on the Lax-Friedrichs scheme described below.

It is important to point out that far more sophisticated schemes exist than the ones presented here. In the applications of these schemes to hyperbolic problems and shock dynamics, high order resolution schemes are often necessary [34], because of the differentiation of the numerical flux function g. However, in our case, because we are solving $\phi_t + H(u) = 0$ rather than $u_t + [H(u)]_x$, the differentiation is not required. Thus, we have found that for almost all practical purposes, the first and second order schemes presented below are adequate.

Before constructing the general schemes, let's return to the example of the propagating curve. Earlier, we attempted to follow the propagation of a simple corner moving with speed $F = 1$. Our attempts to use a central difference approximation failed. In Figure 14, we show what happens if we use the scheme given in Eqn. 51. The exact answer is shown, together with two simulations. The first uses the entropy-satisfying scheme with only 20 points (Fig.14b), the second (Fig.14c) with 100 points. In the first approximation, the entropy condition is satified, but the corner is somewhat smoothed due to the small number of points used. In the more refined calculation, the corner remains sharp, and the exact solution is very closely approximated.

9 First and Second Order Schemes for Convex Speed Functions

Given a convex speed function F (that is, a speed function F such that the resulting Hamiltonian $H = F|\nabla\phi|$ is convex), we can produce the following schemes, first presented in [88]. Start with the equation

$$\phi_t + H(\phi_x, \phi_y, \phi_z) = 0, \tag{52}$$

and approximate it by

$$\phi_i^{n+1} = \phi_i^n - \Delta t\, g\,\big(\frac{\phi_{ijk}^n - \phi_{i-1,j,k}^n}{\Delta x}, \frac{\phi_{i+1,j,k}^n - \phi_{i,j,k}^n}{\Delta x}, \tag{53}$$

23

Exact Solution
*Fig.*14a

Scheme with 20 *Points*
*Fig.*14b

Scheme with 100 *Points*
*Fig.*14c
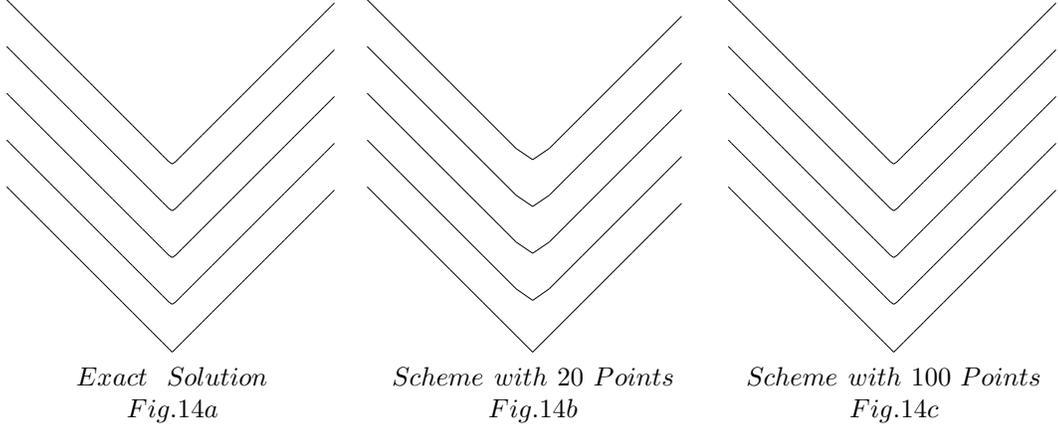
Figure 14: Upwind, Entropy-Satisfying Approximations to the Level Set Equation

$$\frac{\phi_{ijk}^n - \phi_{i,j-1,k}^n}{\Delta y}, \frac{\phi_{i,j+1,k}^n - \phi_{i,j,k}^n}{\Delta y},$$

$$\frac{\phi_{ijk}^n - \phi_{i,j,k-1}^n}{\Delta z}, \frac{\phi_{i,j,k+1}^n - \phi_{i,j,k}^n}{\Delta z}).$$

A multi-dimensional version of this scheme [88] is then given by

$$g_{EO}(u_1, u_2, v_1, v_2, w_1, w_2) = [\max(u_1, 0)^2 + \min(u_2, 0)^2 + \tag{54}$$
$$\max(v_1, 0)^2 + \min(v_2, 0)^2 +$$
$$\max(w_1, 0)^2 + \min(w_2, 0)^2]^{1/2}.$$

Thus we have

1. *First-Order Space Convex:*

$$\phi_{ijk}^{n+1} = \phi_{ijk}^n - \Delta t[\max(F_{ijk}, 0)\nabla^+ + \min(F_{ijk}, 0)\nabla^-] \tag{55}$$

   where

$$\nabla^+ = [\max(D_{ijk}^{-x}, 0)^2 + \min(D_{ijk}^{+x}, 0)^2 + \tag{56}$$
$$\max(D_{ijk}^{-y}, 0)^2 + \min(D_{ijk}^{+y}, 0)^2 +$$
$$\max(D_{ijk}^{-z}, 0)^2 + \min(D_{ijk}^{+z}, 0)^2]^{1/2}$$

$$\nabla^- = [\max(D_{ijk}^{+x}, 0)^2 + \min(D_{ijk}^{-x}, 0)^2 + \tag{57}$$
$$\max(D_{ijk}^{+y}, 0)^2 + \min(D_{ijk}^{-y}, 0)^2 +$$
$$\max(D_{ijk}^{+z}, 0)^2 + \min(D_{ijk}^{-z}, 0)^2]^{1/2}.$$

   Here, we have used a short-hand notation in which $D^{+x}\phi_i^n$ is rewritten as D_i^{+x}, etc.

2. *Second-Order Space Convex:*
   In [88], the above schemes were extended to higher order, using technology from [55]. The basic trick is to build a switch that turns itself off whenever a shock is detected; otherwise, it will use a higher order polynomial approximation of minimal

24

oscillations. These details will not be presented; see [88] for details. The scheme is the same as the above, however this time ∇^+ and ∇^- are given by

$$\nabla^+ = [\,\max(A,0)^2 + \min(B,0)^2 + \tag{58}$$
$$\max(C,0)^2 + \min(D,0)^2 +$$
$$\max(E,0)^2 + \min(F,0)^2]^{1/2}$$

$$\nabla^- = [\,\max(B,0)^2 + \min(A,0)^2 + \tag{59}$$
$$\max(D,0)^2 + \min(C,0)^2 +$$
$$\max(F,0)^2 + \min(E,0)^2]^{1/2}$$

where

$$A = D_{ijk}^{-x} + \frac{\Delta x}{2} m(D_{ijk}^{-x-x}, D_{ijk}^{+x-x}) B = D_{ijk}^{+x} - \frac{\Delta x}{2} m(D_{ijk}^{+x+x}, D_{ijk}^{+x-x}) \tag{60}$$

$$C = D_{ijk}^{-y} + \frac{\Delta y}{2} m(D_{ijk}^{-y-y}, D_{ijk}^{+y-y}) D = D_{ijk}^{+y} - \frac{\Delta y}{2} m(D_{ijk}^{+y+y}, D_{ijk}^{+y-y}) \tag{61}$$

$$E = D_{ijk}^{-z} + \frac{\Delta z}{2} m(D_{ijk}^{-z-z}, D_{ijk}^{+z-z}) F = D_{ijk}^{+z} - \frac{\Delta z}{2} m(D_{ijk}^{+z+z}, D_{ijk}^{+z-z}) \tag{62}$$

and the switch function is given by

$$m(x,y) = \left\{ \begin{array}{cc} \left\{ \begin{array}{cc} x & if \ |x| \le |y| \\ y & if \ |x| > |y| \end{array} \right\} & xy \ge 0 \\ 0 & xy < 0 \end{array} \right\} \tag{63}$$

10 First and Second Order Schemes for Non-Convex Speed Functions

Given a non-convex speed function F (that is, a speed function F such that the resulting Hamiltonian $H = F|\nabla \phi|$ is non-convex), we can follow the work in [89], obtained by replacing the Hamiltonian $F|\nabla \phi|$ with the Lax-Friedrichs numerical flux function produces the following schemes:

1. First-Order Space Non-Convex:

$$\phi_{ijk}^{n+1} = \phi_{ijk}^n - \Delta t [H(\frac{D_{ijk}^{-x} + D_{ijk}^{+x}}{2}, \frac{D_{ijk}^{-y} + D_{ijk}^{+y}}{2}, \frac{D_{ijk}^{-z} + D_{ijk}^{+z}}{2}) \tag{64}$$
$$- \frac{1}{2}\alpha_u(D_{ijk}^{+x} - D_{ijk}^{-x}) - \frac{1}{2}\alpha_v(D_{ijk}^{+y} - D_{ijk}^{-y}) - \frac{1}{2}\alpha_w(D_{ijk}^{+z} - D_{ijk}^{-z})]$$

where α_u (α_v, α_w) is a bound on the partial derivative of the Hamiltonian with respect to the first (second, third) argument, and the non-convex Hamiltonian is a user-defined input function.

2. Second-Order Space Non-Convex:

$$\phi_{ijk}^{n+1} = \phi_{ijk}^n - \Delta t [H(\frac{A+B}{2}, \frac{C+D}{2}, \frac{E+F}{2}) \tag{65}$$
$$- \frac{1}{2}\alpha_u(B-A) - \frac{1}{2}\alpha_v(D-C) - \frac{1}{2}\alpha_w(F-E)]$$

where A, B, C, D, E, and F are defined as above. For details, see [89, 2, 3].

11 Approximations to Curvature and Normals

As discussed above, one of the main advantages of level set formulations is that geometric qualities of the propagating interface, such as curvature and normal direction are easily calculated. For example, consider the case of a curve propagating in the plane. The expression for the curvature of the zero level set assigned to the interface itself (as well as all other level sets) is given by

$$\kappa = \nabla \cdot \frac{\nabla \phi}{|\nabla \phi|} = \frac{\phi_{xx}\phi_y^2 - 2\phi_y\phi_x\phi_{xy} + \phi_{yy}\phi_x^2}{(\phi_x^2 + \phi_y^2)^{3/2}} \tag{66}$$

In the case of a surface propagating in three space dimensions, one has many choices for the curvature of the front, including the mean curvature κ_M and the Gaussian curvature κ_G. Both may be conveniently expressed in terms of the level set function ϕ as

$$\kappa_M = \nabla \cdot \frac{\nabla \phi}{|\nabla \phi|} = \frac{(\phi_{yy} + \phi_{zz})\phi_x^2 + (\phi_{xx} + \phi_{zz})\phi_y^2 + (\phi_{xx} + \phi_{yy})\phi_z^2 - 2\phi_x\phi_y\phi_{xy} - 2\phi_x\phi_z\phi_{xz} - 2\phi_y\phi_z\phi_{yz}}{(\phi_x^2 + \phi_y^2 + \phi_z^2)^{3/2}} \tag{67}$$

$$\kappa_G = \frac{\begin{aligned}&\phi_x^2(\phi_{yy}\phi_{zz} - \phi_{yz}^2) + \phi_y^2(\phi_{xx}\phi_{zz} - \phi_{xz}^2) + \phi_z^2(\phi_{xx}\phi_{yy} - \phi_{xy}^2) \\ &\quad + 2[\phi_x\phi_y(\phi_{xz}\phi_{yz} - \phi_{xy}\phi_{zz}) + \phi_y\phi_z(\phi_{xy}\phi_{xz} - \phi_{yz}\phi_{xx}) \\ &\quad\quad + \phi_x\phi_z(\phi_{xy}\phi_{yz} - \phi_{xz}\phi_{yy})]\end{aligned}}{(\phi_z^2 + \phi_y^2 + \phi_x^2)^2} \tag{68}$$

Construction of the normal itself requires a more sophisticated scheme than simply building the difference approximation to $\nabla \phi$. This is because at corners, the direction of the normal can undergo a jump. This suggests the following technique, introduced in [117]. First, the one-sided difference approximations to the unit normal in each possible direction are formed. All four limiting normals are then averaged to produce the approximate normal at the corner. Thus, the normal n_{ij} is formed by first letting

$$n_{ij}^* \equiv \frac{\phi_x, \phi_y}{(\phi_x^2 + \phi_y^2)^{1/2}} \tag{69}$$

$$= \frac{(D_{ij}^{+x}, D_{ij}^{+y})}{[(D_{ij}^{+x})^2 + (D_{ij}^{+y})^2]^{1/2}} + \frac{(D_{ij}^{-x}, D_{ij}^{+y})}{[(D_{ij}^{-x})^2 + (D_{ij}^{+y})^2]^{1/2}} \tag{70}$$

$$\frac{(D_{ij}^{+x}, D_{ij}^{-y})}{[(D_{ij}^{+x})^2 + (D_{ij}^{-y})^2]^{1/2}} + \frac{(D_{ij}^{-x}, D_{ij}^{-y})}{[(D_{ij}^{-x})^2 + (D_{ij}^{-y})^2]^{1/2}}$$

and then normalizing so that $n_{ij} \equiv n_{ij}^*/|n_{ij}^*|$. If any of the one-sided approximations to $|\nabla \phi|$ is zero, that term is not considered and the weights are adjusted accordingly.

12 Initialization and Boundary Conditions

The time-dependent level set approach requires an initial function $\phi(x, t = 0)$ with the property that the zero level set of that initial function corresponds to the position of the initial front. The original level set algorithm computed the signed distance from each grid point to the initial front which is matched to the zero level set. This is an expensive technique. Many other initial functions are possible, including those that are essentially constant except in a narrow band around the front itself.

The use of a finite computational grid means that we must develop boundary conditions. If the speed function F causes the front to expand (such as in the case $F = 1$), upwind schemes will naturally default to outward-flowing one-sided differences at the boundary of the domain. However, in cases of more complex speed functions, mirror boundary conditions usually suffice.

Part III
Extensions to the Basic Method

13 A Hierarchy of Fast Level Set Methods

The above level set method may be easily programmed. However it is not particularly fast, nor does it make efficient use of computational resources. In this section, we consider a sequence of more sophisticated versions of the basic scheme.

13.1 Parallel Algorithms

The above method updates *all* the level sets, not just the zero level set corresponding to the front itself. The advantage of this approach is that the data structures and operations are extremely clear, and it is a good starting point for building level set codes.

There are a variety of circumstances in which this approach is desirable. If, in fact, all the level sets are themselves important (such as problems encountered in image processing, see a later section), then computation over the entire domain is required. A simple approach is to perform a parallel computation. Since each grid point is updated by a nearest neighbor stencil using only grid points on each side, this technique almost falls under the classification of "embarrassingly parallel". In [114], a parallel version of the level set method was developed for the Connection Machine CM-2 and CM-5. In the CM-2, nodes are arranged in a hypercube fashion, in the CM-5, nodes are arranged in a fat-tree. The code was written in global CMFortran, and at each grid point CSHIFT operators were used to update the level set function. The operation count per time step reduces to $O(1)$, since in most cases the full grid can be placed into physical memory. A time-explicit second order space method was used to update the level set equation.

13.2 Adaptive Mesh Refinement

As a first level of creating a more efficient level set method, an adaptive mesh refinement strategy can be pursued. This is the approach taken by Milne in [82], motivated by the adaptive mesh refinement work in [13]. Adaptivity may be desired in regions where level curves develop high curvature or where speed functions change rapidly. In Figure 15a, we show mesh cells which are hierarchically refined in response to a parent-child relationship around the zero level set of ϕ. Calculations in [82] are performed on both the fine and coarse grid. Grid cell boundaries always lie along coordinate lines and patches do not overlap; in the scheme presented in [82], no attempt is made to align the refined cells with the front.

The data structures for the adaptive mesh refinement are fairly straightforward. However, considerable care must be taken at the interfaces between coarse and fine cells; in particular, the update strategy for ϕ at so-called "hanging nodes" is subtle. These are nodes at the boundary between two levels of refinement which do not have a full set of nearest neighbors required to update ϕ. To illustrate, in Figure 15b, we show a two-dimensional adaptive mesh; the goal is to determine an accurate update strategy for the hanging node marked ○.

The strategy laid out in [82] for updating ϕ at such points is as follows. Consider the archetypical speed function $F(\kappa) = 1 - \epsilon\kappa$.

- The advection term 1 leads to a hyperbolic equation; here, straightforward interpolation of the updated values of ϕ from the coarse cell grid is used to produce the new

Cell Hierarchy *Hanging Nodes*
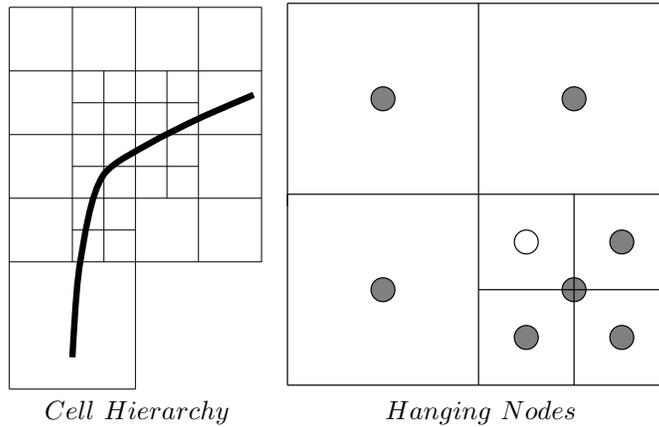
Figure 15: Adaptive Mesh Refinement

value of ϕ at $\circ$. The sophisticated technology in [13] is not required, since we are modeling the update according to the numerical flux function g, not the derivative of the numerical flux function as required for hyperbolic conservation laws.

- In the case of the curvature term $-\epsilon\kappa$, the situation is not as straightforward. The parabolic term cannot be approximated through simple interpolation. In [82], interpolation from updated values on the coarse grid to the fine grid was shown to provide poors answers; if this procedure is employed, the boundary between the two levels of refinement acts as a source of noise, and significant error is generated at the boundary. Instead, values from both the coarse and refined grid next to the hanging node are used to construct a least squares solution for ϕ before the update. This solution surface is then formally differentiated to produce the various first and second derivatives in each component direction. These values are then used to produce the update value for ϕ similar to all other nodes. For details, see [82].
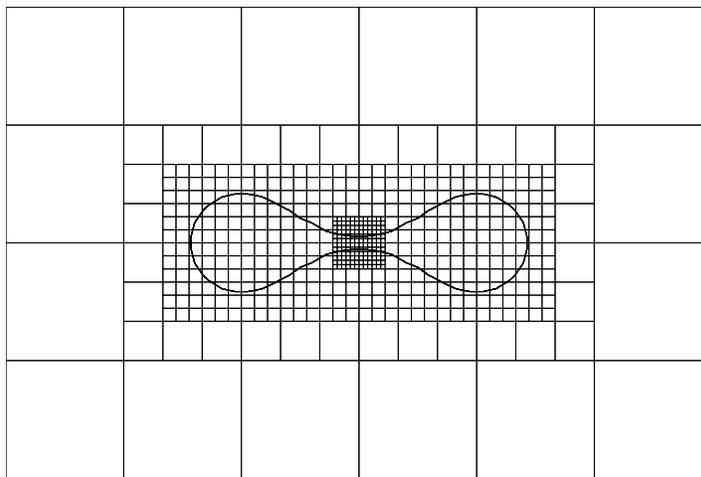
Figure 16: Two-Dimensional Slice of Adaptive Mesh For Propagating Surface

As illustration, in Figure 16, we show a two-dimensional slice of a fully three-dimensional

adaptive mesh calculation of a surface collapsing under its mean curvature. As discussed in a later section, the dumbbell neck pinches off under such a configuration, due to the high positive value of the principle axis of curvature.

13.3   Narrow Banding and Fast Methods

An alternative to the above techniques is particularly valuable when one wants to track a specific front, namely the one associated with the zero level set. There are several disadvantages with the "full-matrix" approach given above.

- *Speed:* Performing calculations over the entire computational domain requires $O(N^2)$ operations in two dimensions, and $O(N^3)$ operations in three dimensions, where N is the number of grid points along a side. As an alternative, an efficient modification is to perform work only in a neighborhood of the zero level set; this is known as the *narrow band approach.* In this case, the operation count in three dimensions drops to $O(kN^2)$, where k is the number of cells in the width of the narrow band. This is a significant cost reduction. Typically, good results can be obtained with about six cells on either side of the zero level.

- *Calculating Extension Variables*: The time-dependent level set approach requires the extension of the speed function F in Eqn. 24 to all of space; this then updates all of the level sets, not simply the zero level set on which the speed function is naturally defined. Recall that three types of arguments may influence the front speed F; local, global, and independent. Some of these variables may have meaning only on the front itself, and it may be both difficult and awkward to design a speed function which extrapolates the velocity away from the zero level set in a smooth fashion. Thus, another advantage of a narrow band approach is that this extension need only be done to points lying in the narrow band.

- The full-matrix approach requires the choice of a time step that applies in response to the maximum velocity over the entire domain, not simply in response to the speed of the front itself. In a narrow band implementation, the time step can be adaptively chosen in response to the maximum velocity field only within the narrow band. This is of significance in problems in which the front speed changes substantially as it moves (for example, due to the local curvature or as determined by the underlying domain). In such problems, the CFL restriction for the velocity field for all the level sets may be much more stringent than the one for those sets within the narrow band.

The above "narrow band method" method was introduced in [28], used in recovering shapes from images in [78], and analyzed extensively by Adalsteinsson in [1].

In Figure 17 we show the placement of a narrow band around an initial front. The entire two-dimensional grid of data is stored in a square array. A one–dimensional object is then used to keep track of the points in this array (dark grid points in Fig. 17 are located in a narrow band around the front of a user-defined width) (see Figure 17). Only the values of ϕ at points within the tube are updated. Values of ϕ at grid points on the boundary of the narrow band are frozen. When the front moves near the edge of the tube boundary, the calculation is stopped, and a new tube is built with the zero level set interface boundary at the center. This rebuilding process is known as "re-initialization".

Thus, the narrow band method consists of the following loop:

- Tag "alive" points in narrow band.
- Build "land mines" to indicate nearby edge.
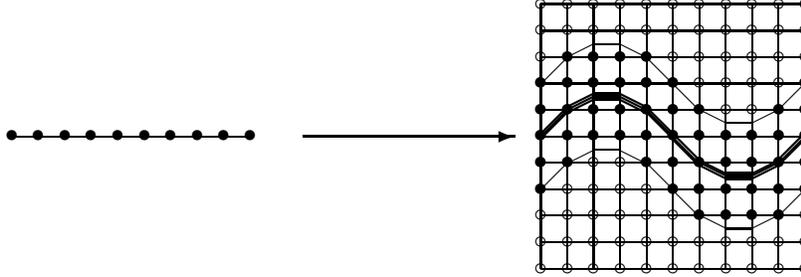
Figure 17: Pointer Array Tags Interior and Boundary Band Points

- Initialize "far away" points outside(inside) narrow band with large positive(negative) values.
- Solve level set equation until a land mine is hit.
- Rebuild, loop.

Use of narrow bands leads to level set front advancement algorithms which are equivalent in complexity to traditional marker methods and cell techniques, while maintaining the features of topological merger, accuracy, and easy extension to multi-dimensions. Typical, the speed up associated with the narrow band method is about ten times faster on a 160×160 grid than the full matrix method. Such a speed up is substantial; in three-dimensional simulations, it can make the difference between computationally intensive problems and those that can be done with relative ease. Details on the accuracy, typical tube sizes, and number of times a tube must be rebuilt may be found in [1].

The above narrow banding technique requires a re-building and re-initialization of a new narrow band around the location of the front. There are several ways to perform this re-initialization, one of which leads to an efficient level set scheme for the particular case of a front propagating with a speed $F = F(x, y)$ where F is of one sign.

13.4 Re-Initialization Techniques: Direct Evaluation, Iteration, Huygens Flowing

13.4.1 Direct Evaluation

A straightforward approach is to find the zero level set by using a contour plotter and then recalculate the signed distance from each grid point to this zero level set to rebuild the band. This technique, first used in [28], can be used to ensure that the level set function stays well-behaved. However, this approach can be expensive, since the front must be explicitly constructed and distances must be calculated to neighboring grid points.

13.4.2 Iteration

An alternative to this was given by Sussman, Smereka, and Osher [122], based on an observation of Morel. Its virtue is that one need not find the zero level set to re-initialize the level set function. Consider the partial differential equation

$$\phi_t = sign(\phi)(1 - |\nabla \phi|) \tag{71}$$

where $sign(\phi)$ gives the sign of ϕ. Given initial data for ϕ, solving the above equation to steady-state provides a new value for ϕ with the property that $—\nabla \phi| = 1$, since

31

convergence occurs when the right-hand-side is zero. The sign function controls the flow of information in the above; if ϕ is negative, information flows one way and if ϕ is positive, then information flows the other way. The next effect is to "straighten-out" the level sets on either side of the zero level set and produce a ϕ function with $|\nabla\phi| = 1$ which in fact corresponds to the signed distance function. Thus, their approach is to periodically stop the level set calculation and solve the above until convergence; if done often enough, the initial data is often close to the signed distance function, and few iterations are required. One disadvantage of this technique is the relative crudeness of the switch function based on checking the sign of the level set equation; considerable motion of the zero level set can occur during the re-initialization, since the sign function does not do an accurate job of using information about the exact location of the front.

13.4.3 Huygens Principle Flowing

An alternative technique, described in [104], is based on the idea of computing crossing times as discussed in [108], and is related to the ideas given by Kimmel in [66]. Consider a particular value for the level set function $\phi_{initial}(x,t)$. With speed function $F = 1$, flow the level set function both forward and backwards in time and calculate crossing times (that is when ϕ changes sign) at each grid point. These crossing times (both positive and negative) are equal to the signed distance function by Huygens' principle. This approach has the advantage that one knows a priori how long to run the problem forward and backward to re-initialize grid points a given distance from the front. This calculation can be performed using a high order scheme to produce accurate values for the crossing times.

This idea of computing crossing times is equivalent to converting the level set evolution problem into a stationary problem. This conversion can be used to develop an extremely fast marching level set scheme for the particular case of solving the level set equation for speed function $F = F(x,y)$, where F is always either positive or negative. This is discussed in the section on fast marching methods for the stationary formulation.

14 Additional Complexities

Since their introduction, the capabilities and applicability of time-dependent level set methods has been considerably refined and extended. In this section, we discuss a few extensions which have proven to be useful in a variety of applications.

14.1 Masking and Sources

Consider the problem of a front propagating with a speed F and subject to the constraint that the evolving interface cannot enter into a region Ω in the domain. This region Ω is referred to as a "mask", since it inhibits all motion. There are several solutions to this problem, depending on the degree of accuracy required.

The simplest solution is to set the speed function F equal to zero for all grid points inside Ω. The location of all points inside Ω can be determined before any calculation is carried out. This technique assures that the front stops within one grid cell of the mask. In Figure 18, we show a plane front propagating upwards with speed $F = 1$ in the upwards direction, with a rectangular block in the center of the domain serving as a mask. In Fig. 18a, the speed function is reset to zero inside the mask region, and as the front propagates upwards it is stopped in the vicinity of the mask and is forced to bend around it.

$Velocity = 0\ Inside\ Rectangle$ $\phi\ Reset\ By\ Mask$
$Fig.18a$ $Fig.18b$
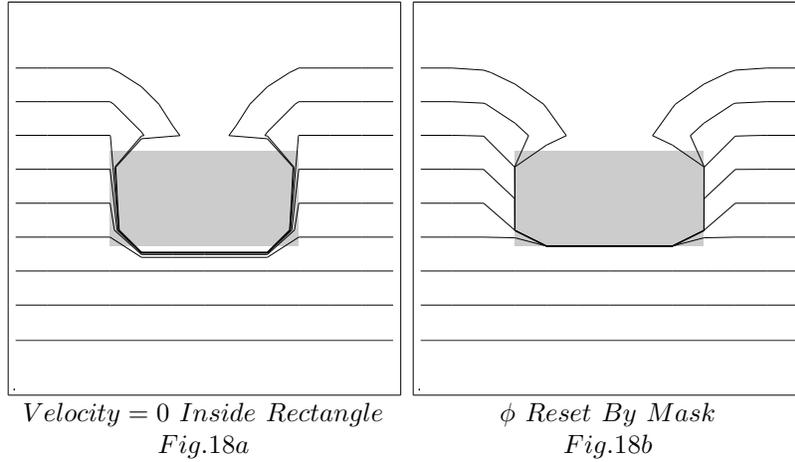
Figure 18: Front Propagating Upwards Around Masking Block: 13×13 Grid

The calculations in Fig. 18a are performed on a very crude 13×13 mesh in order to accentuate a problem with this approach, namely that the front can only be guaranteed to stop within one grid cell of the obstacle itself. This is because the level set method constructs an interpolated speed between grid points, and hence by setting the speed function to zero on and in the mask, the front slows down before it actually reaches the mask. Note that since this means one grid cell *normal* to the mask's boundary, a considerable amount of error can result.

A different fix, which eliminates much of this problem, comes from alternate view. Given a mask area Ω, construct the signed distance function ϕ^Ω by taking the positive distance if inside Ω and the negative distance if outside (note that this is opposite sign choice from the one we typically use). Then we limit motion into the masked region not by modifying the speed function but instead by resetting the evolving level set function. Let $\phi^{(*)}$ be the value produced by advancing the level set ϕ^n one time step. Then let

$$\phi^{n+1} = \max(\phi^*, \phi^\Omega). \tag{72}$$

This resets the level set function so that penetration is not possible; of course, this is only accurate to the order of the grid. Results using this scheme are shown in Figs. 18b. Again, we have used a very coarse grid to accentuate the differences.

If we consider now the opposite problem, in which a region Ω acts as a source, the solution is equally straightforward, and given by $\phi^{n+1} = \min(\phi^*, -\phi^\Omega)$; this is the technique used in [95], inspired by the work in [91].

14.2   Discontinuous Speed Functions and Sub-grid Resolution

Let us generalize the above problem somewhat, and imagine that we want to solve a interface propagation problem in which there is a clear discontinuous speed function. For example, one may want to track the propagation of an interface through materials which correspond to differing media, across which propagation rates change quite sharply. As an example, consider again the evolution of the upwards propagating front, but this time the rectangular block slows the speed to $1/2$ (that is, $F = 1$ outside Ω and $F = .5$ inside and on Ω. The standard level set will interpolate between these two speeds, and the results obtained will depend on the placement of the underlying grid; substantial variation in result will occur depending on whether a grid line lies directly on, below, or above the bottom edge of the rectangular block.
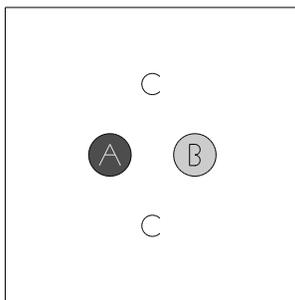
Figure 19: Regions A and B Expand into Region C

In order to accurately solve this problem, we need some sub-grid information about the speed function to correctly construct the speed function for those cells that lie only partially within Ω. Such a technique can be devised, motivated by the idea of the volume-of-fluid methods discussed earlier. Given a region Ω, before any calculation proceeds we construct the cell fraction Vol_{ij}^{Ω}, which is a number between 0 and 1 for those cells that have at least one grid point in Ω and one outside Ω. This cell fraction corresponds to the amount of Ω material in the cell. These values are stored, and a list is kept of such boundary cells. We then proceed with the level set calculation, letting F be given by its value in the corresponding region. However, we modify the speed function for those cells that are marked as boundary cells. At the beginning of the time step, compute the volume fraction Vol_{ij}^{ϕ} for the zero level set in each cell; this may be approximately done without explicitly finding the zero level set through a least squares fit. This value is then compared with the stored value Vol_{ij}^{Ω} to provide an appropriate speed.

14.3   Multiple Interfaces

As initially designed in [88], the level set technique applies to problems in which there is a clear distinction between an "inside" and "outside". This is because the interface is assigned the zero level value between the two regions. Extensions to multiple (more that two) interfaces have been made in some specific cases. In the case in which interfaces are passively transported and behave nicely, one may be able to use only one level set function and judiciously assign different values at the interfaces. For example, the zero level set may correspond to the boundary between two regions A and B, with the level set value 10 corresponding to the interface between two regions B and C. If A and C never touch, then this technique may be used to follow the interfaces in some cases.

However, in the more general case evolving the emergence and motion of triple points, a different approach is required, since many different situations can occur (see, for example, [17, 123]). Consider the following canonical example, as illustrated in Figure 19. Regions A and B are both circular disks growing into region C with speed unity in the direction normal to the interface. At some point, the interfaces will touch and meet at a triple point, where a clear notion of "inside" and "outside" cannot be assigned in a consistent manner.

A level set approach to this problem has been proposed by Bence, Merriman and Osher [81]; they move each interface separately for one time step, find the interfaces of the various fronts, and then rebuild level set functions. This technique requires re-initializing the pairwise level set functions; any of the techniques described earlier can be used. Before describing this technique, we discuss a later set of techniques presented in [109, 104], applicable to many cases, and which do not require any such re-initialization.

Influence Matrix	on A	on B	on C
Effect of A	1.0	0.0	1.0
Effect of B	0.0	1.0	1.0
Effect of C	0.0	0.0	1.0
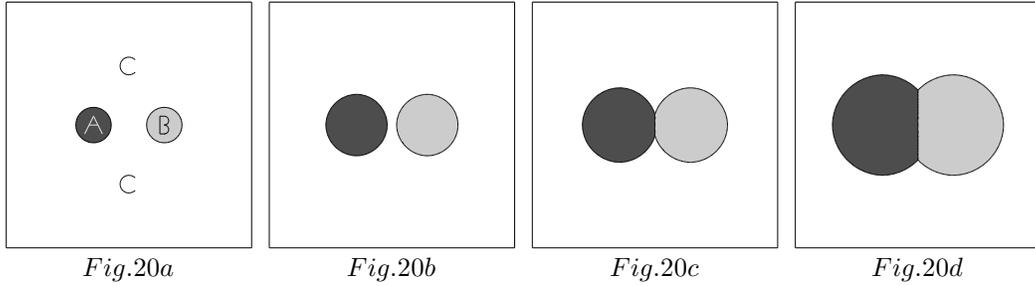
Figure 20: A and B Move into C with speed 1, Stop at each other

We will then return to the algorithm presented in [81] in the case of motion driven by surface tension where some sort of re-initialization is required.

The key idea in each method lies in recasting the interface motion as the motion of one level set function for each material. In some sense, this is what was done in the re-ignition idea given in [95], where the front was a flame which propagated downstream under a fluid flow. This front was re-ignited at each time step at a flame holder point by taking the minimum of the advancing flame and its original configuration around the flame holder, thus ensuring that the maximum burned fluid is achieved.

In general, imagine N separate regions and a full set of all possible pairwise speed functions F_{IJ} which describe the propagation speed of region I into region J: F is taken as zero if Region I cannot penetrate J. The idea is to advance each interface to obtain a trial value for each interface with respect to motion into every other region, and then combine the trial values in such a way as to obtain the maximum possible motion of the interface.

In general then, proceed as follows. Given a Region I, obtain $N-1$ trial level set functions ϕ_{IJ}^* by moving the Region I into each possible Region J, J=1,N ($J \neq I$) with speed F_{IJ}. During the motion of Region I into Region J, assume that all other regions are impenetrable, that is, use the masking rule given by Eqn. 72. We then test the penetrability of the Region J itself, leaving the value of ϕ_{IJ}^* unchanged if $F_{IJ} \neq 0$, else modifying it with the maximum of itself and $-\phi_{JI}^*$. Finally, to allow Region I to evolve as much as possible, we take the minimum over all possible motions to obtain the new position; this is the re-ignition idea described earlier. Complete details of the approach may be found in [104].

Three examples, taken from [104], are shown to illustrate this approach. Given regions A, B, and C, the *influence matrix* describes the interaction of the various regions with each other. The interaction of each region with itself is 1, hence the matrix has ones on the diagonal. The interaction of any pair of regions is required to be zero in one of the two interactions.

In Figure 20, regions A and B expand with unit speed into Region C, but cannot penetrate each other. They advance and meet; the boundary between the two becomes a vertical straight line.

Next, we consider a problem with different evolution rates. In Figure 21, region A grows with speed 1 into region C (and region C grows with speed 0 into region A), and region B grows with speed 2 into region C. Once they come into contact, region

35

| Fig.21a | Fig.21b | Fig.21c | Fig.21d |
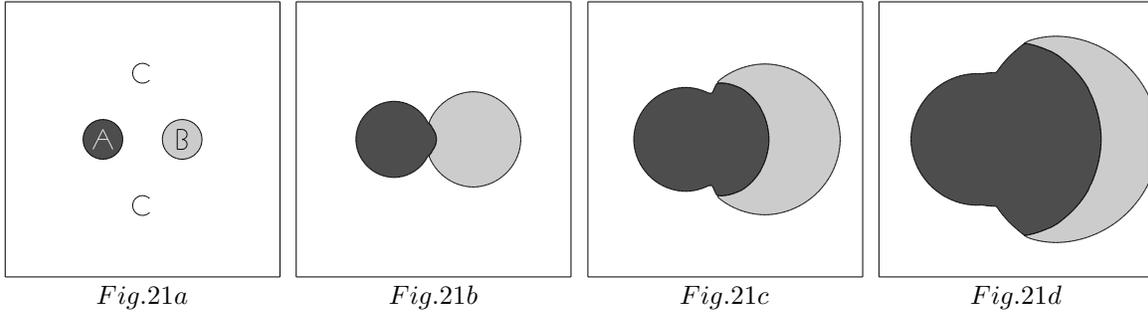
Influence Matrix	on A	on B	on C
Effect of A	1.0	3.0	1.0
Effect of B	0.0	1.0	2.0
Effect of C	0.0	0.0	1.0

Figure 21: A into C with speed 1, A into B with speed 3, B into C with speed 2

| Fig.22a T = 0.0 | Fig.22b T = 2.0 | Fig.22c T = 3.0 | Fig.22c T = 5.0 |
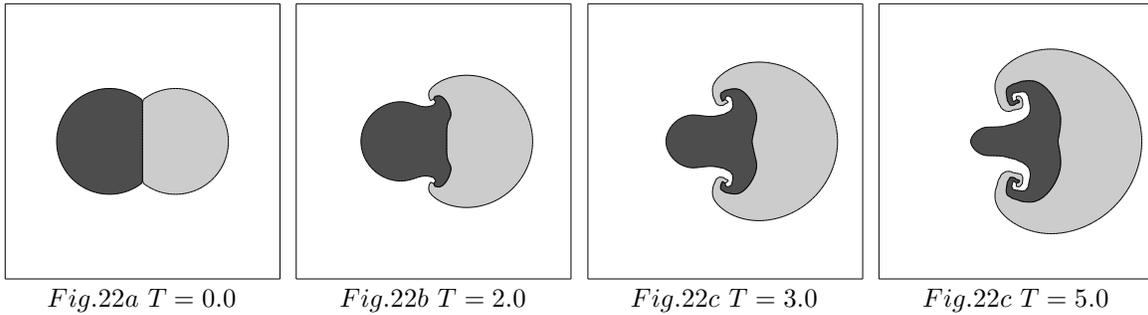
Influence Matrix	on A	on B	on C
Effect of A	1.0	1.0	0.0
Effect of B	0.0	1.0	1.0
Effect of C	1.0	0.0	1.0

Figure 22: Spiraling Triple Point: 98x98 Grid

A dominates region B with speed 3, thus region B grows through Figure 21c, and then is "eaten up" by the advancing region A. Note what happens: region A advances with speed 3 to the edge of region B, which is only advancing with speed 1 into region C. However, region A cannot pass region B, because *its* speed into region C is slower than that of region B.

Finally, in Figure 22 the motion of a triple point between regions A, B, and C is shown. Assume that region A penetrates B with speed 1, B penetrates C with speed 1, and C penetrates A with speed 1. The exact solution is given by a spiral with no limiting tangent angle as the triple point is approached. The triple point does not move; instead, the regions spiral around it. In Figure 22, results are shown from a calculation on a 98×98 grid. Starting from the initial configuration, the regions spiral around each other, with the leading tip of each spiral controlled by the grid size. In other words, we are unable to resolve spirals tighter than the grid size, and hence that controls the fine scale description of the motion. However, we note that the triple point remains fixed. A series of additional calculations using this approach may be found in [109].

14.4 Triple Points

Consider now the case of a triple point motion in which the speed of each interface is driven by curvature, which may correspond to surface tension. Imagine a triple point, in which each of the three regions is attempting to move according to their own curvature. In Figure 23a, we show an initial configuration on the left and a final state in Figure 23b in the center, which consists of the three lines meeting in equal angles of 120 degrees.

*Fig.23a : Initial State   Fig.23b : Final State   Fig.23c : Level Set*
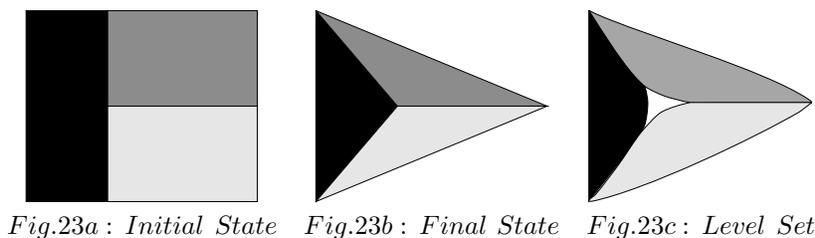
Figure 23: Evolution of Triple Point under Curvature

If one attempts to apply the level set method for multiple interfaces described in the previous sections, a difficulty occurs because each level function attempts to move away from the others, creating a gap, In Fig. 23c, we show this gap developing when a level set technique is applied to the final state.

Two different level set type algorithms were introduced in [81] for tackling this problem. The first can be viewed a "fix" to the above problem; the second is a wholly different level set approach.

To begin, the problem with the above calculation is that the various level sets pull apart. In [81], the level set functions are reset every time step to hold the triple point in place, that is,

$$\phi_i = \phi_i - \max \phi_j \tag{73}$$

where the maximum is taken over all $j \neq i$. This keeps the triple point in place, however, as explained in [81], the cost is that the level set functions can develop spontaneous zero crossings later in time. A remedy is to re-initialize all the level sets using any of the re-initialization techniques described in the previous section. With those two added steps in the algorithm, level set methods can easily handled some interesting problems concerning triple points (see [81]).

A considerably different approach was also presented in [81]. The "DGCDM" algorithm works by applying applying a reaction-diffusion type equation to a characteristic function assigned to each region, which is one inside the region and zero outside. This algorithm works by exploiting the link between curvature flow and a diffusion equation, along the lines of the material discussed earlier. The basic idea is that a diffusion term is applied, and then a sharpening term is executed which sharpens up the solution. The net effect is to evolve the boundary line under curvature. This is a very clever algorithm, and can be applied to multiple interfaces. A series of fascinating calculations are presented in [81], for additional work on this topic, see [17].

14.5 Building Extension Velocity Fields

What happens when the speed of the moving front has meaning only on the front itself? This is a common occurrence in areas such as combustion, material science, and fluid mechanics, where the philosophy of embedding the front as the zero level set of a family

of contours can be problematic. In fact, the most difficult part of level set methods is this "extension" problem, and will be a central focus of the Part IV.

Recall the division of arguments in the speed function F given in Eqn. 1. Front-based arguments are those that depend on geometric quantities of the front, such as curvature and normal vector, and have a clear meaning for all the level sets. Independent variables are equally straightforward, since their contribution to the speed makes no reference to particular information from the front itself.

The troublesome variables are the so-called "global" variables, which can arise from solving differential equations on either side of the interface. Briefly, there are at least four ways to extend a velocity from the front to the grid points.

1. At each grid point, find the closest point on the front. This was the technique used in [78], and may be done efficiently in many cases by tracing backwards along the gradient given by $\nabla \phi$.

2. Evaluate the speed function off the front using an equation that only has meaning on the front itself. This is the technique used in the crystal growth/dendritic solidification calculations employed in [117], where a boundary integral is evaluated both on and off the front.

3. Develop an evaluation technique that assigns artificial speeds to the level set going through any particular grid point. For example, in the etching/deposition simulations of [2, 3, 4], visibility of the zero level set must be evaluated away from the front itself.

4. Smearing the influence of the front. In the combustion calculations of [95] and the fluid dynamics calculations of [122], the influence of the front is mollified to neighboring grid points on which an appropriate equation is solved.

Part IV

Approximating the Stationary Level Set Formulation: A New Fast Marching Level Set Method

Consider the special case of a monotonically advancing front, that is, a front moving with speed F where F is always postive (or negative). Previously, we have produced a stationary level set equation, namely

$$|\nabla T|F = 1. \tag{74}$$

which is simply a static Hamilton-Jacobi equation; if F is only a function of position, this becomes the well-known Eikonal equation. A large body of research has been devoted to studying these types of equations, we refer the interested reader to [10, 71, 118], to name just a few. At the same time, there are a wide collection of schemes to solve this problem, see, for example, [11, 46, 96, 87]. Here, we introduce an entirely new and extremely fast method for solving this equation. It relies on a marriage between our narrow band technique and a fast heapsort algorithm, and can be viewed as an extreme one-cell version of our narrow band technique.

For ease of discussion, we limit ourselves to a two-dimensional problem inside a square from $[0, 1] \times [0, 1]$ and imagine that the initial front is along the line $y = 0$; furthermore, we assume that we are given a positive speed function $F(x, y)$ that is periodic in x. Thus, the front propagates upwards off the initial line, and the speed does not depend on the orientation of the front (it depends only on *independent variables*, using our earlier terminology). Using our approximation to the gradient, we are then looking for a solution in the unit box to the equation

$$F_{ij} \left[\max(D_{ij}^{-x}T, 0)^2 + \min(D_{ij}^{+x}T, 0)^2 + \max(D_{ij}^{-y}T, 0)^2 + \min(D_{ij}^{+y}T, 0)^2) \right] = 1 \tag{75}$$

where $T(x, 0) = 0$.

Since Eqn. 75 is in essence a quadratic equation for the value at each grid point (assuming the others are held fixed), we can iterate until convergence by solving the equation at each grid point, selecting the largest possible value as the solution in accordance with obtaining the correct viscosity solution. An iterative algorithm for computing the solution to this problem was introduced by Rouy and Tourin in [96]. Typically, one iterates several times through the entire set of grid points until a converged solution is reached.

15   A fast marching level set method

The key to constructing a fast marching algorithm is the observation that the upwind difference structure of equation 75 means that information propagates "one way", that is, from smaller values of T to larger values. Hence, our algorithm rests on "solving" equation 75 by building the solution outwards from the smallest time value T. Our idea is to sweep the front ahead in an upwind fashion by considering a set of points in narrow band around the existing front, and to march this narrow band forward, freezing the values of existing points and bringing new ones into the narrow band structure. The key is in the selection of *which* grid point in the narrow band to update. The technique is
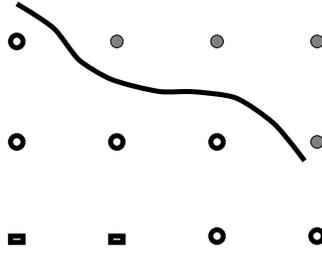
Figure 24: Narrow band approach to marching level set method

easiest to explain algorithmically, see Figure 24. We imagine that we want to propagate a front through an N by N grid with speed F_{ij} giving the speed in the normal direction at each grid point. Here the set of grid points $j = 1$ correspond to the y axis, and we assume that $F_{ij} > 0$.

Algorithm:

1. Initialize
   (a) (Alive Points: Grey Disks): Let A be the set of all grid points $\{i, j = 1\}$; set $T_{i,1} = 0.0$ for all points in A.
   (b) (Narrow Band Points: Black Circles): Let $NarrowBand$ be the set of all grid points $\{i, j = 2\}$, set $T_{i,1} = dy/F_{ij}$ for all points in $NarrowBand$.
   (c) (Far Away Points: Black Rectangles): Let $FarAway$ be the set of all grid points $\{i, j > 2\}$, set $T_{i,j} = \infty$ for all points in $FarAway$.

2. Marching Forwards
   (a) Begin Loop: Let $(i_{\min}, j_{\min})$ be the point in $NarrowBand$ with the smallest value for T.
   (b) Add the point $(i_{\min}, j_{\min})$ to A; remove it from $NarrowBand$.
   (c) Add to the narrow band list any neighboring points $(i_{\min} - 1, j_{\min})$, $(i_{\min} + 1, j_{\min})$, $(i_{\min}, j_{\min} - 1)$, $(i_{\min}, j_{\min} + 1)$ that are not $Alive$. If the neighbor is in $FarAway$, remove it from that list.
   (d) Recompute the values of T at all neighbors according to Eqn. 75, selecting the largest possible solution to the quadratic equation.
   (e) Return to top of Loop:

We take periodic boundary conditions where required. Assuming for the moment that it takes no work to determine the member of the narrow band with the smallest value of T, the total work required to compute the solution at all grid points is $O(N^2)$, where calculation is performed on an N by N grid.

Why does the above algorithm work? Since we are always locating the smallest value in the narrow band, its value for T must be correct; other narrow band points or far away points with larger T values cannot affect it. The process of recomputing the T values at neighboring points (that have not been previously accepted) cannot yield a value smaller than any of that at any of the accepted points, since the correct viscosity solution is obtained by selecting the *largest* possible solution to the quadratic equation. Thus, we can march the solution outwards, always selecting the narrow band grid point with minimum trial value for T, and readjusting neighbors. Another way to look at this is that each minimum trial value begins an application of Huygen's principle, and the expanding wave front touches and updates all others.

Figure 25: Matrix of neighboring values

15.1 Proof that the algorithm constructs a viable solution

Here, we prove that the above algorithm produces a solution which everywhere satisfies the discrete version of our equation, which is given by

$$\left[\max(\max(D_{ij}^{-x}T, 0), -\min(D_{ij}^{+x}T, 0))^2 + \max(\max(D_{ij}^{-y}T, 0), -\min(D_{ij}^{+y}T, 0))^2\right] = f_{ij}^2 \tag{76}$$

where $f_{ij}^2 = 1/F_{ij}^2$. We shall give a constructive proof. Since the values of $T(x, y, z)$ are built by marching forwards from the smallest value to the largest, we need only show that whenever a "trial" value is converted into an "alive" value, none of the recomputed neighbors obtain new values less than the accepted value. If this is true, then we will always be marching ahead in time, and the thus the correct "upwind" nature of the differencing will be respected. We shall prove our result in two dimensions; the three-dimensional proof is the same.

Thus, consider the matrix of grid values given in Figure 25. Our argument will follow the computation of the new value of T in the center grid point to replace the value of ?, based on the neighboring values. We will assume, without loss of generality, that the value A at the left grid point is t he smallest of all "trial" values, and prove that when we recompute the value at the center grid point (called $T_{recomputed-from-A}$), it cannot be less than A. This will prove that the upwinding is respected, and that we need not go back and readjust previously set values. We shall consider the four cases that (1) none of the neighbors B, C, or D, are "alive", (2) one of these neighbors is "alive", (3) two of the neighbors are "alive", and (4) all three of these neighbors are "alive". [2]

A, B, C and D are "trial", A is the smallest

In this case, all of the neighbors around the center grid point are either "trial" or set to $FarAway$. Since A is the smallest such value, we convert that value to "alive" and recompute the value at the center grid point. We now show that the recomputed value $A \leq T_{recomputed-from-A} \leq A + f$.

1. Suppose $A + f \leq \min(B, D)$. Then $T_{recomputed-from-A} = (A + f)$ is a solution to the problem, since only the difference operator to the left grid point is non-zero. footnoteWe are absorbing the grid size Δx into the inverse speed function f.

2. Suppose $A + f \geq \min(B, D)$. Then, without loss of generality, assume that $B \leq D$. We can solve the quadratic equation

$$(T_{recomputed-from-A} - A)^2 + (T_{recomputed-from-A} - B)^2 = f^2 \tag{77}$$

The discriminant is non-negative when $f \geq \frac{(B-A)}{\sqrt{2}}$, which must be true since we assumed that $A + f \geq B$ and hence $f \geq (B - A)$. Thus, a solution exists, and it is easy to check that this solution must then be greater than or equal to B and thus falls into the required range. Furthermore, we see that $T \leq A + f$, since the second term on the left is non-negative.

---

[2] Recall that "alive" means that their T values are less than A. Here, we are using the notation that the symbol A stands for both the grid point and its T value.

Thus, we have shown that $A \leq T_{recomputed-from-A} \leq A+f$, and therefore $T_{recomputed-from-A}$ cannot be less than the just converted value A.

This case will act a template for the other cases.

B is "alive", A, C and D are "trial", A is the smallest of the trial values

In this case, A has just been converted, since it is the smallest of the trial values. We shall prove when we recalculate $T_{recomputed_A}$, its new value must still be greater than A. At some previous stage, when B was converted from trial to alive, the values of A, C and D were all trial values, and hence must have been larger. Then this means that when B was converted from trial to alive, we had the previous case above, and hence $B \leq T_{recomputed-from-B} \leq B + f$; furthermore, since the value at the center was *not* chosen as the smallest trial value, we must have that $A \leq B + f$. By the above case, we then have that $B \leq A \leq T_{recomputed-from-A} \leq B + f$, and hence the recomputed value cannot be less than the just converted value of A.

C is "alive", A, B and D are "trial", A is the smallest of the trial values

In this case, due to the direction of the upwind differencing, the value at C is the contributor in the x direction, the acceptance of A does not affect the recomputation, and the case defaults into the first case above.

The remaining cases are all the same, since the differencing takes the smallest values in each coordinate direction. The proof in three dimensions is identical.

15.2   Finding the smallest value

The key to an efficient version of the above technique lies in a fast way of locating the grid point in the narrow band with the smallest value for T. We use a variation on a heapsort algorithm, see Press et. al. [92] and Sedgewick [103], with the additional feature of back pointers. In more detail, imagine that the list of narrow band points is initially sorted in a heapsort so that the smallest member can be easily located. We store the values of these points in the heapsort, together with their indices which give their location in the grid structure. We keep a companion array which points from the two-dimensional grid to the location of that grid point in the heapsort array. Finding the smallest value is easy. In order to find the neighbors of that point, we use the pointers from the grid array to the heapsort structure. The values of the neighbors are then recomputed, and then the results are bubbled upwards in the heapsort until they reach their correct locations, at the same time readjusting the pointers in the grid array. This results in an $O(logN)$ algorithm for the total amount of work, where N is the number of points in the narrow band. For implementation details and further application of this technique, see [111, 110, 116].

The above technique considered a flat initial interface for which trial values at the narrow band points could be easily initialized. Suppose we are given an arbitrary closed curve or surface as the initial location of the front. In this case, we use the original narrow band level set method to initialize the problem. First, label all grid points as "far away" and assign them T values of ∞. Then, ina very small neighborhood around the interface,construct the signed distance function from the initial hypersurface Γ. Propagate that surface both forwards and backwards in time until a layer of grid points is crossed in each direction, computing the signed crossing times as in [108]. Then collect the points with negative crossing times as "alive" points with T value equal to the crossing time, and the points with positive crossing times as narrow band points with T value equal to the positive crossing times. Then begin the fast marching algorithm.

16 Other Speed Functions: When does this method work?

Our fast marching method as designed applied to speed functions F which depend only on position. In such cases, other forms for the gradient approximation can be used; for example, Rouy and Tourin [96], namely

$$F_{ij} \left[\max(\max(D_{ij}^{-x}T, 0), -\min(D_{ij}^{+x}T, 0))^2 + \max(\max(D_{ij}^{-y}T, 0), -\min(D_{ij}^{+y}T, 0))^2 \right] = 1 \tag{78}$$

How general is our new technique? Suppose now we consider the more general case of stationary level set equation:

$$|\nabla T| F = 1. \tag{79}$$

We begin be rewriting this in the standard form of a static Hamiltonian, namely

$$H(T_x, T_y, T_z) = 1. \tag{80}$$

We already have a scheme for the case where $H = \nabla T$. Some variations on this Hamiltonians important in computer vision, such as

$$H = \max(|T_x|, |T_y|, |T_z|) \qquad H = |T_x| + |T_y| + |T_z| \tag{81}$$

may be approximated in a straightforward manner using the any of the above entropy-satisfying approximations to the individual gradients. Our fast marching method will work in these cases. Wnen the speed F depends in a subtle way on the value of ∇T (for example, in some problems in etching and deposition discussed in a later section, the situation is more delicate.

When will the technique work? Here, we present an intuitive perspective; complete details may be found in [110, 116]. Suppose H is convex and always positive (or always negative), and suppose the approximation to H satisfies two properties:

- First, the approximation scheme is consistent
- Second, at each grid point the scheme only makes use of smaller neighboring values when updating the value at that point (this is the upwindness requirement), and cannot produce a new value which is less than any of the neighbors

Then we can expect that our upwind sweeping method will work; searching for the smallest trial value will provide a consistent way of sweeping through the mesh and constructing the soluton surface T. Complete details and many other schemes may be found in [110, 116].

17 Some Clarifying Comments

The time-dependent level set method and the stationary level set method each require careful construction of upwind, entropy-satisfying schemes, and make use of the dynamics and geometry of front propagation analyzed in [106]. However, we note that the time-dependent level set method advances the front **simultaneously**, while the stationary method constructs "scaffolding" to build the time solution surface T one grid point at a time. This means that the time at which the surface crosses a grid point (that is, its T value) may be found before other positions of that front at that time are determined.

As such, there is **no** notion of a time step in the stationary method; one is simply constructing the stationary surface in an upwind fashion.

This means that if one is attempting to solve a problem in which the the speed of a front depends on the current position of the front (such as in the case of visibility), or on subtle orientations in the front (such as in sputter yield problems), it is not clear how to use the stationary method, since the front is being constructed one grid point at a time.

The stationary method works because we were able to construct a simple approximation to the gradient. This was possible because the speed function F did **not** depend on the orientation of the front, nor on issues like visibility. Thus, returning to our earlier categorization of speed functions, our fast scheme works in cases where the speed F only depends on independent variables, such as in the case of photolithography development. Upwind entropy-satisfying schemes which can be transported to this fast stationary scheme for the case of more general speed functions F are more problemmatic, and discussed in detail in [110].

To summarize,

- The stationary method is convenient for problems in which the front speed depends on independent variables, such as a photoresist rate function, and only applies if the speed function does not change sign.

- The time-dependent level set method is designed for more delicate speed functions, and can accurately fronts evolving under highly complex arguments

In the next part, we discuss a variety of applications which employ both the time-dependent level set method and the fast marching method for monotonically advancing fronts.

Part V
Applications

In this part, we present a series of applications of both the time-dependent level set method and the fast marching level set method to propagating interfaces. This is only a subset of the level set applications in the literature; throughout, we provide references to many other applications.

18   Geometry

In this section, we consider application of level set methods to problems in the geometric evolution of curves and surfaces. The motion will depend solely on local geometric properties such as normal direction and curvature; nonetheless, this is a fascinating and rich area.

18.1   Curvature Flow

Suppose we are given a hypersurface in R^n propagating with some speed $F(\kappa)$. Previously, we have considered speed functions of the form $F(\kappa) = 1 - \epsilon\kappa$, where κ is the curvature. Let us now focus on a special speed function, namely $F = -\kappa$, where κ is the curvature. This corresponds to a geometric version of the heat equation; large oscillations are immediately smoothed out, and long-term solutions correspond to dissipation of all information about the initial state. As we shall see in later sections, curvature motion plays an important role in many applications such as a modeling term for surface tension in flexible membranes and a viscous term in physical phenomena.

   The remarkable work of Gage and Grayson investigated the motion of a simple closed curve collapsing under its curvature. First, Gage [48, 49] showed that any convex curve moving under such a motion remains convex and must shrink to a point. Grayson [52] followed this work with a stunning proof that *all* curves must shrink to a round point, regardless of their initial shape.

   In Figure 26, we take an odd-shaped initial curve and view this as the zero level set of a function defined in all of R^2. Here, for illustration, we have ϕ such that $\phi < 0$ as black and $\phi > 0$ as white, thus the zero level set is the boundary between the two. As the level curves flow under curvature, the ensuing motion carries each to a point which then disappears. In the evolution of the front, one clearly sees that the large oscillations disappear quickly, and then as the front becomes circular, motion slows, and the front eventually disappears.
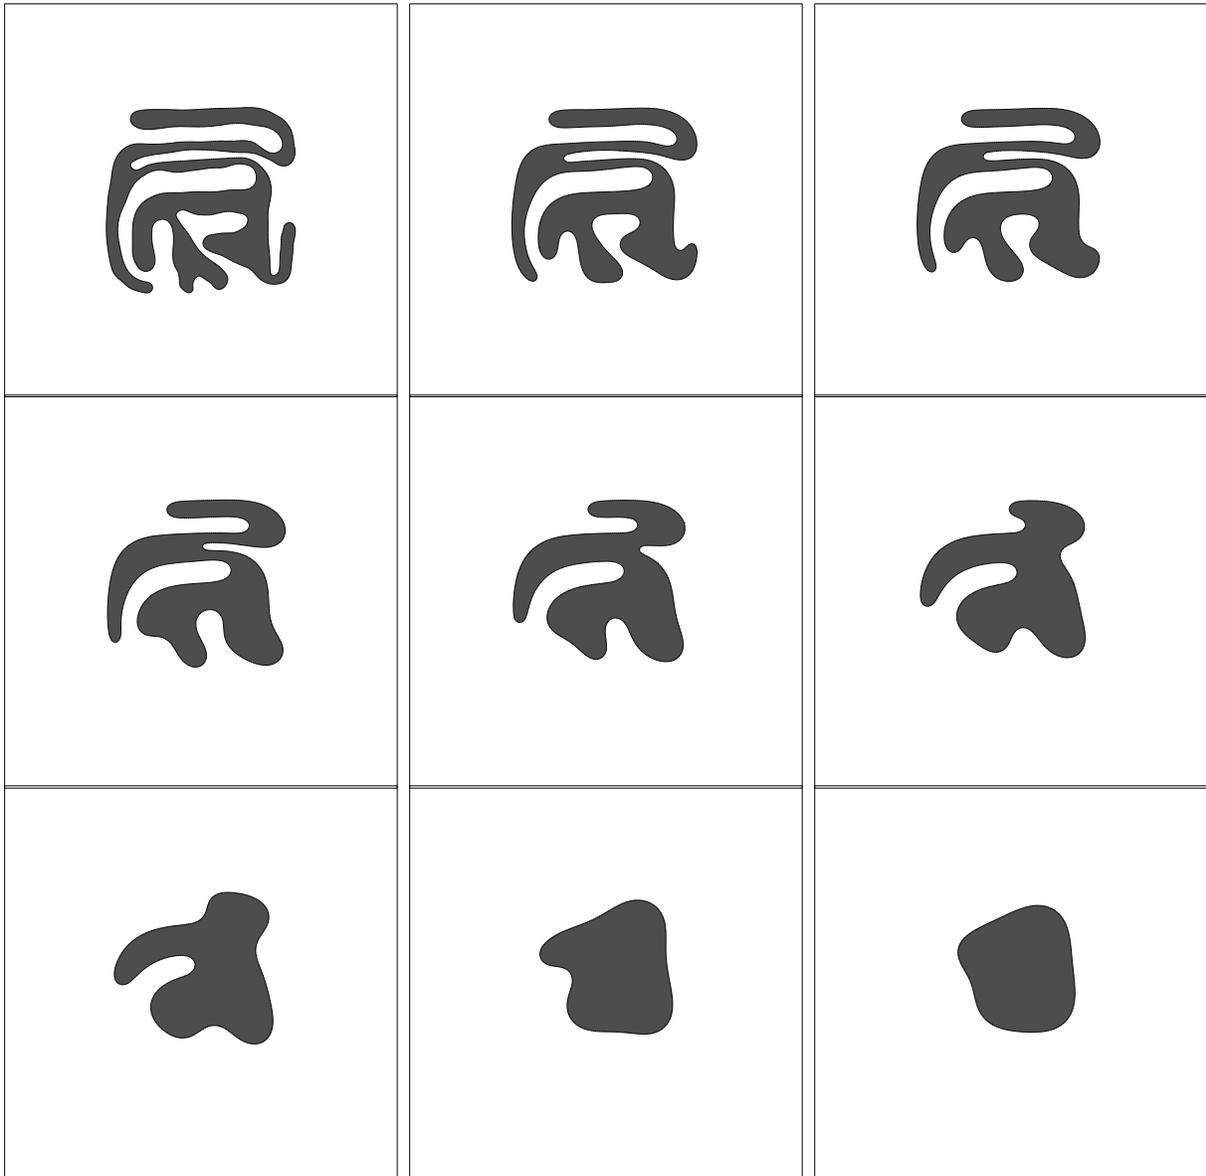
Figure 26: $F(\kappa) = -\kappa$:
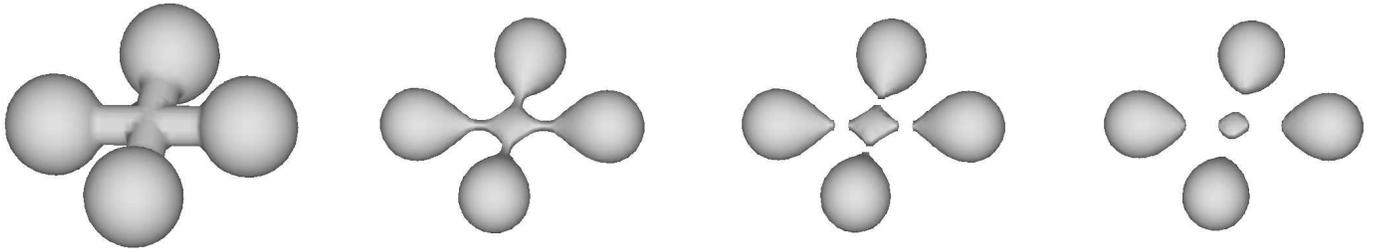
Figure 27: Collapse of Two-handled Dumbbell

In three dimensions, flow under mean curvature does not necessarily result in a collapse to a sphere. Huisken [58] showed that convex shapes shrink to spheres as they collapse, analogous to the result of Gage. However, Grayson showed that non-convex shapes may in fact *not* shrink to a sphere [53], and provided the counterexample of the dumbbell. A narrow handle of a dumbbell may have such a high inner radius that the mean curvature of the saddle point at the neck may still be positive, and hence the neck will pinch off.

As illustration, in Figure 27, taken from [29], we show two connected dumbbells collapsing under mean curvature. As the intersection point collapses, the necks break off and leave a remaining "pillow"-like region behind. This pillow region collapses as well, and eventually all five regions disappear.

Finally, what about self-similar shapes? In two dimensions, it is clear that a circle collapsing under its own curvature remains a circle, this can be seen by integrating the ordinary differential equation for the changing radius. In three dimensions, a sphere is self-similar under mean curvature flow, since its curvature is always constant. Angenent [7] proved the existence a self-similar torus which preserves the balance between the competing pulls towards a ring and a sphere.

In order to devise an algorithm to produce self-similar shapes, two things are required. First, since hypersurfaces get smaller as they move under their curvature, a mechanism is needed to "rescale" their motion so that the evolution can be continued towards a possible self-similar shape. And second, a way of pushing the evolving fronts back towards self-similarity is required. Chopp has accomplished both in a clever numerical algorithm that produces a family of self-similar surfaces, see [27]. His family comes from taking regular polyhedra (for example, a cube), and drilling holes in each face. The resulting figure then evolves according to auxiliary level set equation which contains the re-scaling as part of the equation of motion. Two such self-similar surfaces are shown in Figure 28.

$Self - Similar\ \ Cube\ \ with\ \ Holes\quad Self - Similar\ \ Octahedron\ \ with\ \ Holes$

$Fig.\ 28a \qquad\qquad\qquad Fig.\ 28b$

Figure 28: Self-Similar Shapes

18.2 Grid Generation

Imagine that one is given a closed body, either as a curve in two space dimensions, or a surface in three space dimensions. In many situations, one wishes to generate a logically rectangular, body-fitted grid around or inside this body. By logically rectangular, we mean that each node of the grid has four neighbors (in two dimensions; in three dimensions, there are 6 neighbors); by body-fitted, we mean that the grid aligns itself with the body so that one set of coordinate lines matches the body itself. This grid generation problem is difficult in part because of the competing desires of uniformity in cell area and mesh orthogonality; we refer the interested to [67].

Level set techniques offer an interesting technique for generating such grids. The idea, as presented in [108], is to exploit the geometric nature of the problem and view the body itself as the initial position of an interface which must be advanced outwards away from the body. The initial position of the interface and its position at later times forms one set of grid lines; its orthogonal set forms the other. The body is propagated outwards with speed $F = 1 - \epsilon\kappa$; by finding the zero level set at discrete times, the set of coordinate lines that encircle the body is found. Construction of transverse lines normal to the body are obtained by following trajectories of $\nabla\phi$. Additional node adjustment is possible through application of additional smoothing operators. Node placement on the boundary and the ensuing exterior/interior grid can be automatically controlled. For details, see [108].

This technique roughly falls into the category of a hyperbolic solver. However, by solving the correct evolution equation for an advancing front, we avoid the difficulties of shock formation and colliding characteristics that plague most hyperbolic techniques. User intervention is kept to a minimum; for the most part, grids are generated automatically without the need to adjust parameters.

In Figure 29, taken from [108], we show a variety of grids constructed using this level set approach, starting with relatively smooth grids and ending with a three-dimensional grid around an indented dumbbell. As can be seen, interior and exterior grids can be created, with the capability of handling significant corners and cusps. The grids are automatically created; there has been no adjustment of parameters in the creation of these different grids.

49

Figure 29: Body-Fitted Grids Generated Using Level Set Approach

Figure 30: An Image $I(x, y)$ with Noise

18.3 Image Enhancement and Noise Removal

The previous sections concerned geometrical motion of a particular hypersurface of interest. Next, we turn to a level set problem in which *all* the level sets have meaning, and must be evolved.

The goal in this section is to apply some of the level set methodology to image enhancement and noise removal. To do so, we first need a few definitions. Define an *image* to be an intensity map $I(x, y)$ given at each point of a two-dimensional domain. The range of the function $I(x, y)$ depends on the type of image; for black and white image the range is either 0 or 255, for grey-scale images $I(x, y)$ is a function mapped between 0 and 255, and, for color images, $I(x, y)$ is a vector-valued function into some color-space, typically either RGB or HSI.

Given such an image, two goals are to remove noise from the image (see Figure 30) without sacrificing useful detail and to enhance or highlight certain features. A straightforward and widely used approach is the Gaussian filter, in which both 1-D and 2-D signals are smoothed by convolving them with a Gaussian kernel; the degree of blurring is controlled by the characteristic width of the Gaussian filter. If the image is viewed as a surface, this convolution will reduce spikes as they blend into the background values. In this sense, the Gaussian removes noise. However, the Gaussian is an isotropic operator; it smoothes in all directions, and sharp boundaries will also be blurred. The goal is to improve upon this basic idea and remove noise without being forced into too much blurring. A variety of techniques have been introduced to improve upon this basic idea, including anisotropic diffusion schemes which perform intraregion smoothing in preference to interregion smoothing, see Perona and Malik [90], as well as Wiener filters, and wavelet schemes.

Alvarez, Lions and Morel ([5]) introduced a significant advancement in noise removal by employing, in part, some of the above ideas about curvature flow and level set equations. Consider the equation

$$ I_t = F|\nabla I| \qquad where \qquad F = \nabla \cdot \frac{\nabla I}{|\nabla I|} = \kappa \tag{82} $$

where again $\nabla \cdot \frac{\nabla I}{|\nabla I|} = \kappa$. This is our standard curvature evolution equation. An attractive quality of this motion is that sharp boundaries are preserved; smoothing takes place inside a region, but not across region boundaries. Of course, as shown by Grayson's theorem, eventually all information is removed as each contour shrinks to zero and disappears.

An alternative approach is due to Rudin, Osher and Fatemi ([97]). They take a total variation approach to the problem which leads to a level set methodology and a very similar curvature-based speed function, namely $F(\kappa) = \kappa/|I|$; see also [86]. Following these works, a variation on these two approaches was produced by Sapiro and Tannenbaum ([100]). In that work, a speed function of the form $F(\kappa) = \kappa^{1/3}$ was employed. In each of these schemes, all information is eventually removed through continued application of the scheme. Thus, a "stopping criteria" is required.

A recent level set scheme for noise removal and image enchancement was introduced by Malladi and Sethian in [75, 74]. The scheme results from returning to the original ideas of curvature flow, and exploiting a "min/max" function which correctly selects the optimal motion to remove noise. It has two highly desirable features:

1. There is an intrinsic, adjustable definition of scale within the algorithm, such that all noise below that level is removed, and all features above that level are preserved.

2. The algorithm stops automatically once the sub-scale noise is removed; continued application of the scheme produces no change.

These two features are quite powerful, and lead to a series of open questions about the morphology of shape and asymptotics of scale-removal; for details, see [74].

18.3.1   The Min/Max Flow

Consider the equation

$$\phi_t = \bar{F}|\nabla\phi|. \tag{83}$$

A curve collapsing under its curvature will correspond to speed $\bar{F} = \kappa$. Now, consider two variations on the basic curvature flow, namely

- $\bar{F}(\kappa) = \min(\kappa, 0.0)$
- $\bar{F}(\kappa) = \max(\kappa, 0.0)$

Here, we have chosen the negative of the signed distance in the interior, and the positive sign in the exterior region. The effect of flow under $\bar{F}(\kappa) = \min(\kappa, 0.0)$ is allow the inward concave fingers to grow outwards, while suppressing the motion of the outward convex regions. Thus, the motion halts as soon as the convex hull is obtained. Conversely, the effect of flow under $\bar{F}(\kappa) = \max(\kappa, 0.0)$ is to allow the outward regions to grow inwards while suppressing the motion of the inward concave regions. However, once the shape becomes fully convex, the curvature is always positive and the flow becomes the same as regular curvature flow.

Our goal is to select the correct choice of flow that smoothes out small oscillations, but maintains the essential properties of the shape. In order to do so, we discuss the idea of the min/max switch.

Consider the following speed function, introduced in [75] and refined considerably in [74]:

$$\bar{F}^{Stencil=k}_{\min/\max} = \begin{cases} \min(\kappa, 0) & \text{if } Ave^{R=kh}_{\phi(x,y)} < 0 \\ \max(\kappa, 0) & \text{if } Ave^{R=kh}_{\phi(x,y)} \geq 0 \end{cases} \tag{84}$$

where $Ave^{R=kh}_{\phi(x,y)}$ is defined as the average value of ϕ in a disk of radius $R = kh$ centered around the point (x, y). Here, h is the step size of the grid. Thus, given a "StencilRadius" k, the above yields a speed function which depends on the value of ϕ at the point (x, y), the average value of ϕ in neighborhood of a given size, and the value of the curvature of the level curve going through (x, y).

Initial Boundary
"Noisy'' Shape

Min/Max Flow :
Stencil Radius = 0; ($T = \infty$)

Min/Max Flow :
Stencil Radius = 1; ($T = \infty$)

Continued Min/MaxFlow :
*Stencil Radius* = 2; ($T = \infty$)
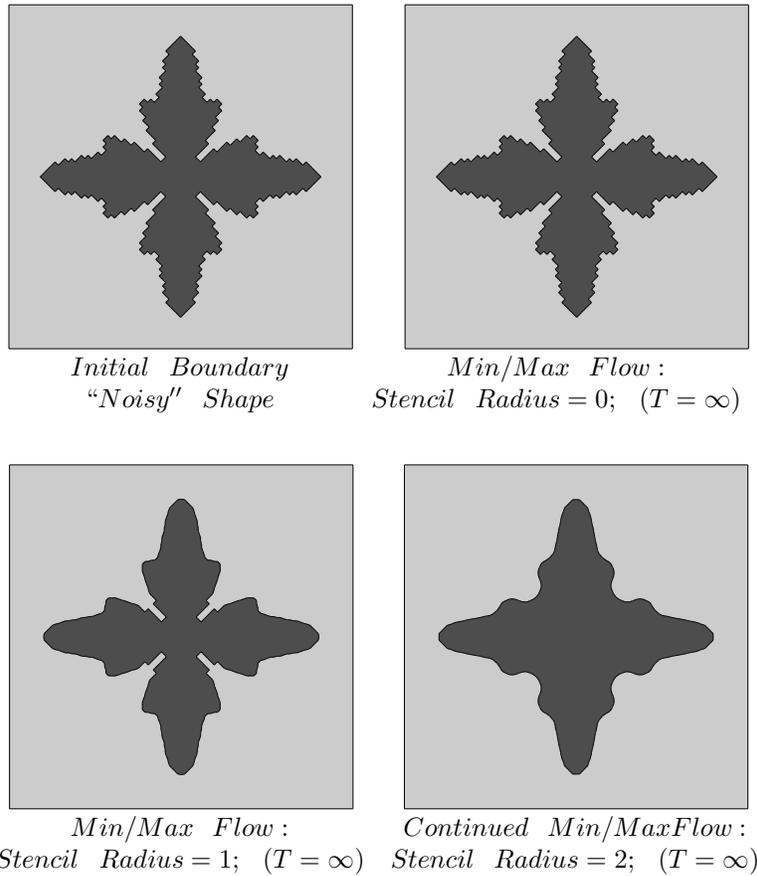
Figure 31: Motion of Star-Shaped Region with Noise under Min/Max Flow at Various Stencil Levels

We can examine this speed function in some detail. For ease of exposition, consider a black region on a white background, chosen so that the interior has a negative value of ϕ and the exterior a positive value of ϕ.

- StencilRadius $k = 0$

  If the radius $R = 0$ ($k = 0$), then choice of $\min(\kappa, 0)$ or $\max(\kappa, 0)$ depends only the value of ϕ. All the level curves in the black region will attempt to form their convex hull, when seen from the black side, and all the level curves in the white region will attempt to form <u>their</u> convex hull. The net effect will be no motion of the zero level set itself, and the boundary will not move.

- StencilRadius $k = 1$

  If the average is taken over a stencil of radius kh, then some movement of the zero level corresponding to the boundary is possible. If there are some oscillations in the front boundary on the order of one or two pixels, then the average value of ϕ at the point (x, y) can have a different sign than the value at (x, y) itself. In this case, the flow will act as if it were selected from the "other side", and some motion will be allowed until these first-order oscillations are removed, and a balance between the two sides is agained reached. Once this balanced is reached, further motion is suppressed.

- <u>StencilRadius k</u>

  By taking averages over a larger stencil, larger amounts of smoothing are applied to the boundary. In other words, decisions about where features belong are based on larger and larger perspectives. Once features on the order of size k are removed from the boundary, balance is reached and the flow stops automatically. As an example, let $k = \infty$. Since the average will compute to a value close to the background color, on this scale all structures are insignificant, and the max flow will be chosen everywhere, forcing the boundary to disappear.

To show the results of this hierarchical flows, we start with an initial shape in Figure 31a and first perform the min/max flow under steady-state is reached with stencil size zero in Figure 31b. In this case, the steady-state is achieved right away, and the final state is the same as the initial state. Min/max flow is then performed until steady-state is achieved with stencil size $k = 1$ in Figure 31c, and then min/max flow is again applied with a larger stencil until steady-state is achieved in Figure 31d.

These results can be summarized as follows:

- The min/max flow switch selects the correct motion to diffuse the small-scale pixel notches into the boundary.

- The larger, global properties of the shape are maintained.

- Furthermore, and equally importantly, the flow stops once these notches are diffused into the main structure.

- Edge definition is maintained, and, in some global sense, the area inside the boundary is preserved.

- The noise removal capabilities of the min/max flow are scale-dependent, and can be hierarchically adjusted.

- The scheme requires only a nearest neighbor stencil evaluation.

18.3.2   Extension of Min/Max Scheme to Grey-Scale, Texture, and Color Images

The above technique applies to black and white images. An extension to grey-scale images can be made by replacing the fixed threshold test value of 0 with a value that depends on the local neighborhood. As designed in [75], let $T_{threshold}$ be the average value of the intensity obtained in the direction perpendicular to the gradient direction. Note that since the direction perpendicular to the gradient is tangent to the iso-intensity contour through (x, y), the two points used to compute are either in the same region, or the point (x, y) is an inflection point, in which the curvature is in fact zero and the min/max flow will always yield zero. By choosing a larger stencil we mean computing this tangential average over endpoints located further apart.

Formally then, our min/max scheme becomes:

$$\bar{F}_{min/max} = \begin{cases} \max(\kappa, 0) & \text{if } Average(x, y) < T_{threshold} \\ \min(\kappa, 0) & \text{otherwise} \end{cases} \qquad (85)$$

Further details about this scheme may be found in [74]. In that work, these techniques are applied to a wide range of images, including salt-and-pepper, multiplicative and Gaussian noise applied to black and white, grey scale, textured, and color images.

| 50.0% *Noise* | *Restored* | 80.0% *Noise* | *Restored* |
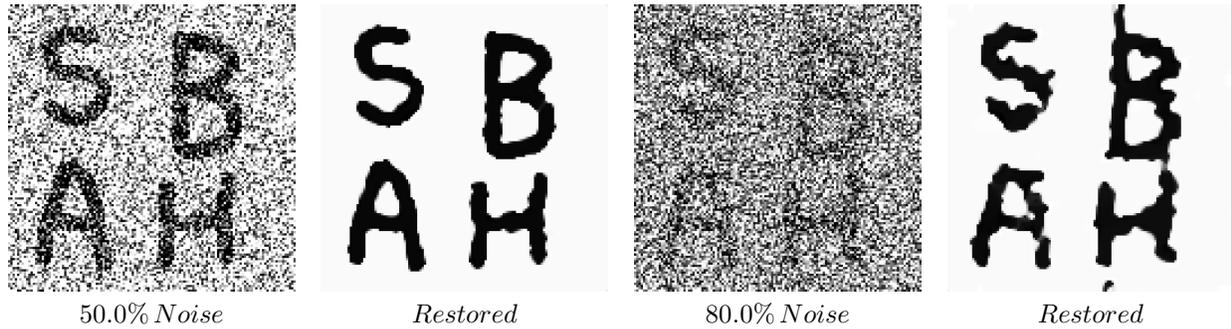
Figure 32: Image restoration of Binary Images with Grey-Scale Salt-and-Pepper Noise Using Min/Max Flow: Restored shapes are final shape obtained ($T = \infty$).


(a) 40% Noise   (b) Min/Max Flow   (c) Cont.: Larger Stencil
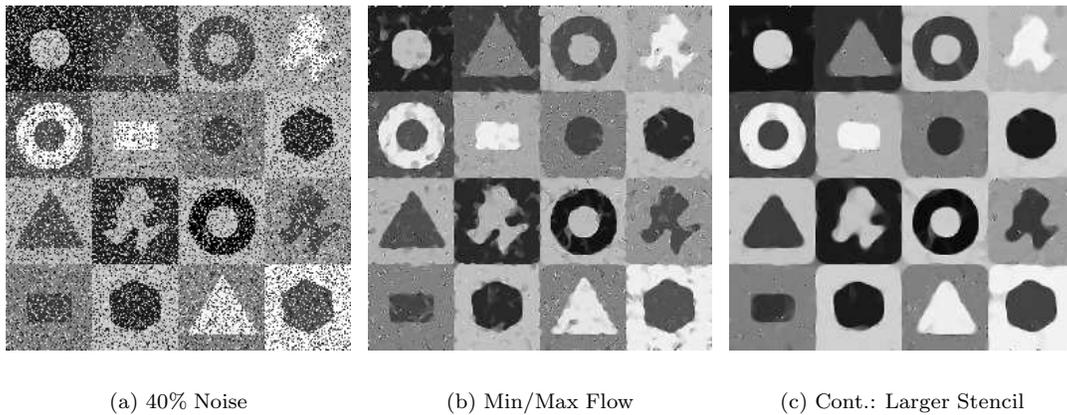
Figure 33: Min/Max Flow Applied to Grey-Scale Image

18.3.3   Results

In this section, we provide a few examples of this min/max flow. In Figure 32, 50% and 80% grey-scale noise is added to a black and white image of a hand-written character. The noise is added as follows: $X\%$ noise means that at $X\%$ of the pixels, the given value is replaced with a number chosen with uniform distribution between 0 and 255. Here, the min/max switch function is taken relative to the value 127.5 rather than zero. The restored figures are converged. Continued application of the scheme yields almost no change in the results.

Next, salt-and-pepper gray-scale noise is removed from a grey-scale image. The results are obtained as follows. Begin with 40% noise in Figure 33a. First, the min/max flow from Eqn. 85 is applied until a steady-state is reached (Figure 33b. This removes most of the noise. The scheme is then continued with a larger threshold stencil for the threshold to remove further noise (Figure 33c). For the larger stencil, we compute the average $Average(x, y)$ over a larger disk, and compute the threshold value $T_{threshold}$ using a correspondingly longer tangent vector.

| *Original* | *Reconstructed* |

Figure 34: Continuous Gaussian Noise added to Image

Finally, we show the effects of this scheme applied to an image upon which 100% Gaussian grey-scale noise has been superimposed a random component drawn from a Gaussian distribution with mean zero is added to each (every) pixel. Figure 34 shows the "noisy" original together with the reconstructed min/max flow image.

19 Combustion, Crystal Growth, and Two-Fluid Flow

In this section, we consider some applications of the level set methodology to a large and challenging class of interface problems. These problems are characterized by physical phenomena in which the front acts as a boundary condition to a partial differential equation, and the solution of this equation controls the motion of the front. In combustion problems, the interface is a flame, and both exothermic expansion along the front and flame-induced vorticity drive the underlying fluid mechanics. In crystal growth and dendritic formation, the interface is the solid/liquid boundary, and is driven by a jump condition related to heat release along the interface. In two fluid problems, the interface represents the boundary between two immiscible fluids of significantly different densities and/or viscosities, and the surface tension along this interface plays a significant role in the motion of the fluids.

From an algorithmic perspective, the significant level set issue is that information about the speed of the front itself must be somehow transferred to the Eulerian framework that updates the level set function at the fixed grid points. This is significant challenge for two reasons:

- In many situations, the interface velocity is determined by the interaction of local geometric quantities of the front itself (such as curvature) with global variables on either side of the interface (for example, jumps in velocity, heat, or concentration of species). If these global variables are calculated on a grid, it may be difficult to extend the values to the front itself where they are required to evaluate the speed

function F. However, these quantities are only known at grid points, not at the front itself, where the relationship has meaning.

- It may be very difficult to extend this interface velocity back to the grid points (that is, to the other level sets). This problem, known as the *extension problem*, see [2, 3], must be solved in order for the level set method to work; some mechanism of updating the grid points in the neighborhood of the zero level set is required.

In this section, we discuss three applications areas where these problems have been solved. The common link in these sections is the presence of a term in the equation represented by a Dirac delta function along the interface. The interested reader is referred to the literature cited below for a detailed description of the algorithms and results.

## 19.1 Turbulent Combustion of Flames and Vorticity, Exothermicity, Flame Stretch and Wrinkling

In [95], a flame is viewed as an infinitely thin reaction zone, separating two regions of different but constant densities. The hydrodynamic flow field is two-dimensional and inviscid, and the Mach number is vanishingly small. This corresponds to the equations of zero Mach number combustion, introduced in [72]. The flame propagates into the unburnt gas at a prescribed flame speed S_u which depends on the local curvature, due to the focusing/de-focusing of heating effects as a function of flame shape.

As the reactants are converted to products (that is, as the material makes the transition from "unburnt" to "burnt"), the local fluid undergoes volume increase known as exothermic expansion associated with the density jump. At the same time, pressure gradients tangential to the flame cause different accelerations in the light and heavy gases. This causes a production of vorticity (known as baroclinic torque) across the flame, since the pressure gradient is not always aligned with the density gradient. Together, the burning of the flame acts as source of vorticity and volume for the underlying hydrodynamic field, both of which in turn affect the evolution of the flame interface.

This model presents a significant challenge for a level set method. The flame is tracked by identifying the flame interface as the zero level set of the level set function. The curvature is determined using the expression given in Eqn. 66. The vortical field is represented by a collection of vortex blobs as in vortex method, see [32, 107]. The exothermic field is determined by solving a Poisson's equation on the underlying grid with right-hand-side given by smearing the Dirac delta function to the neighboring grid points. The no-flow boundary is satisfied by the addition of a potential flow that exactly cancels the existing flow field. Finally, the tangential stretch component is evaluated by tracing the values of the tangential velocity from the given position backwards along the normal to the front to evaluate the change in tangential velocity. For complete details, see [95].

Figure 35, taken from [95], shows two results from this algorithm. We compare an anchored flame, with upstream turbulence imposed by a statistical distribution of positive and negative vortices. The goal here is to understand the effect of exothermicity and flame-induced vorticity on the flame wrinkling and stability. In Fig. 35a, we show an anchored flame in the oncoming turbulent field; here, different time snapshots are superimposed upon each other to show the flame "brush". In Fig. 35b, we turn on the effects of both volume expansion (a large density jump) and vorticity generation along the flame front. The resulting flow field generates a significantly wider flame brush, as the vorticity induces flame wrinkling and the exothermicity affects the surrounding flow field.

In the above application of the level method, the front acted as source of volume and
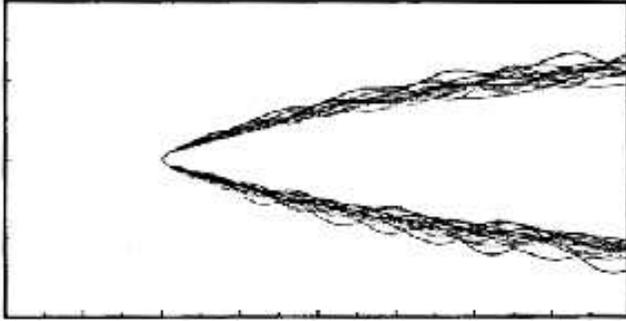
Flame with Exothermicity
*Fig.*35*a*
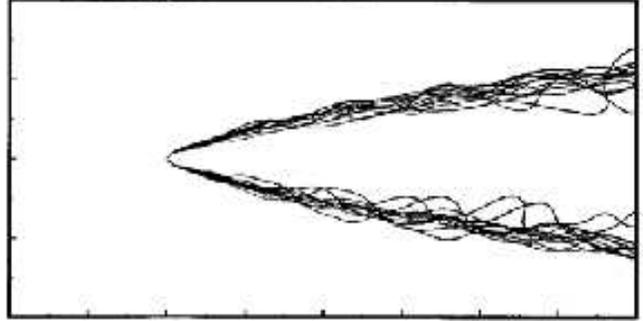
Flame with Exothermicity and Vorticity
*Fig.*35*b*

Figure 35: Comparison of Flame Brush

vorticity. In the next application, developed in [117], the interface motion is controlled by a complex jump condition.

19.2 Crystal Growth and Dendritic Solidification

Imagine a container filled with a liquid such as water, which has been so smoothly and uniformly cooled below its freezing point that the liquid does not freeze. The system is now in a "metastable" state, where a small disturbance such as dropping a tiny seed of the solid phase into the liquid will initiate a rapid and unstable process known as **dendritic solidification**. The solid phase will grow from the seed by sending out branching fingers into the distant cooler liquid nearer the undercooled wall. This growth process is *unstable* in the sense that small perturbations of the initial data can produce large changes in the time-dependent solid/liquid boundary.

Mathematically, this phenomenon can be modeled by a moving boundary problem. The temperature field satisfies a heat equation in each phase, coupled through two boundary conditions on the unknown moving solid/liquid boundary, as well as initial and boundary conditions. First, the normal velocity of the interface depends on the jump in the normal component of the heat flux across the interface. Second, the temperature at the interface itself depends on the local curvature and the velocity. Thus, the goal is to incorporate the influence of the front on the heat solvers across the interface. For further details, see [19, 84].

A variety of techniques are possible to numerically approximate these equations of motion. One approach is to solve the heat equation in each phase and try to move the boundary so that the two boundary conditions are satisfied. Another approach is to recast the equations of motion as a single integral equation on the moving boundary and solve the integral equation numerically. References for these various techniques are given in [117].

In [117] a hybrid level set/boundary integral approach was developed, which includes the effects of undercooling, crystalline anisotropy, surface tension, molecular kinetics, and initial conditions. The central idea is to exploit a transformation due to Strain [119] which converts the equations of motion into a single, history-dependent boundary integral equation on the solid/liquid boundary, given by

$$\epsilon_\kappa(n)\kappa + \epsilon_V(n)V + U + H \int_0^t \int_{\Gamma(t')} K(x, x', t - t') \, V(x', t') dx' dt' = 0 \qquad (86)$$
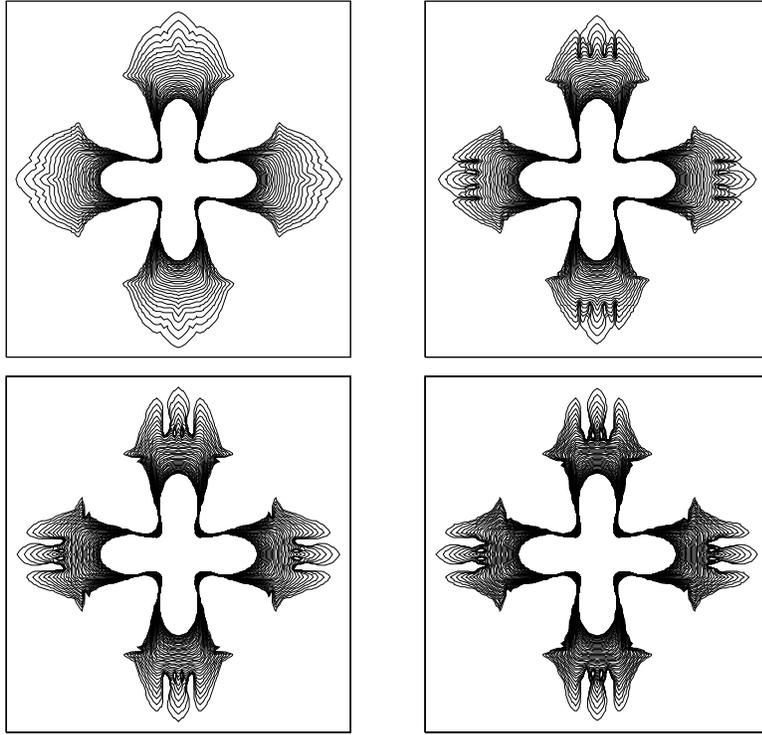
Upper Left: $H = .75$ Upper Right: $H = .833$
Lower Left: $H = .916$ Lower Right: $H = 1.0$

Figure 36: Effect of Changing Latent Heat

for all x on the interface $\Gamma(t)$. Here, K is the heat kernel, ϵ_κ and ϵ_V are constants, U is the temperature, V is the normal velocity of the interface, and H is the latent heat of solidification. Note that the velocity V depends not only on the position of the front but also on its previous history. Thus, as shown in [119], information about the temperature off the front is stored in the previous history of the boundary. This can be evaluated by a combination of fast techniques, see [119, 120, 121, 54].

In Figure 36, we show one example from [117] in which the effect of changing the latent heat of solidification H is analyzed. Since the latent heat controls the balance between the pure geometric effects and the solution of the history-dependent heat integral, increasing H puts more emphasis on the heat equation/jump conditions. Calculations are performed on a unit box, with a constant undercooling on the side walls of $u_B = -1$. The kinetic coefficient is $\epsilon_V = .001$, the surface tension coefficient is $\epsilon_\kappa = .001$, there is no crystalline anisotropy. The initial shape was a perturbed circle. A 96×96 mesh is used with time step $\Delta t = .00125$. The calculations are all plotted at the same time.

In the calculations shown, H is varied smoothly. In Fig. 36a, $H = .75$; the dominance of geometric motion serves to create a rapidly evolving boundary that is mostly smooth. H is increased in each successive figure, ending with $H = 1.0$ in Fig. 36d. As the latent heat of solidification is increased, the growing limbs expand outwards less smoothly, and instead develop flat ends. These flat ends are unstable and serve as precursors to tip splitting. We also note that the influence of the heat integral slows down the evolving boundary, as witnessed by the fact that all the plots are given at the same time. Presumably increasing latent heat decreases the most unstable wavelength, as described by linear stability theory. The final shape shows side-branching, tip-splitting, and the strong effects of the side walls.

20   Two-Phase Flow

In this last section, we discuss level set methods applied to problems of two-phase flow.

Two early applications of fluid dynamical problems using level set methods to track the interface are the projection method calculations of compressible gas dynamics of Mulder, Osher and Sethian [83] and the combustion calculations of Zhu and Sethian [128]. Each viewed the interface as the zero level set, and tracked this interface as a method of separating the two regions.

To begin, in [83] the evolution of rising bubbles in compressible gas dynamics was studied. The level set equation for the evolving interface separating two fluids of differing densities was incorporated inside the conservation equations for the fluid dynamics. Both the Kelvin-Helmholtz instability and the Rayleigh-Taylor instability were studied; the density ratio was about 30 to 4, and both gases were treated as perfect gases. Considerable discussion was devoted to the advantages and disadvantages of embedding the level set equation as an additional conservation law.

Next, in the combustion calculations presented in [128], the interface represented a flame propagating from the burnt region into the unburnt region. Unlike the above calculations concerning flame stability in flame holders, in these calculations the flame was viewed as a "cold flame"; that is, the hydrodynamic flow field affected the position of the flame, but the advancing flame did not in turn affect the hydrodynamic field. In these calculations, the hydrodynamic field was computed using Chorin's projection method (see [31]), coupled to the level set approach. The problem under study was the evolution of a flame inside a swirling two-dimensional chamber; and the results showed the intermixing that can occur, the creating of pockets of unburnt fuel surrounded by burnt pockets.

These two works were followed by the work of Chen, Hou, Merriman, Osher, Smereka, and Sussman in [25, 122] using projection methods coupled to the level set equation to study the motion of incompressible, immiscible fluids where steep gradients in density and viscosity exist across the interface, and the role of surface tension is crucial.

There are three key aspects of these fascinating calculations. First, using a formulation first developed by Brackbill, Kothe, and Zemach [14], the surface tension generated by the level set curvature expression can be smoothed using a mollified delta function to the neighboring grid points; this smoothing denotes a "thickness" for the interface layer, and allows the role of surface tension to be transported to the grid for inclusion in the projection method. Second, the contribution due to surface tension is converted from a delta function to the Heaviside formulation and incorporated as such into the projection step. This eliminates the standard numerical instabilities and oscillations that plague attempts to directly difference the delta function itself. Third, the level set lines cease to correspond to the distance function due to the significant variation in fluid velocities across the interface; to redistribute the level set contours, the re-initialization idea using iteration on the sign of ϕ described in Eqn. 71 was developed in [122].

The calculations performed using these techniques show a wide range of applications concerning falling drops, colliding drops, and the role of surface tension, and open the door to a range of important fluid dynamics applications. We refer the reader to [122, 25] for further details.

Figure 37: Minimal Surface: Catenoid

21 Constrained Problems: Minimal Surfaces and Shape Recovery

Another set of applications include problems in which the motion of the front is constrained by external boundary conditions. In this section, we briefly review some work on the construction of minimal surfaces and shape recovery from images.

21.1 Minimal Surfaces

The basic problem may be stated as follows: Consider a closed curve $\Gamma(s)$ in R^3. The goal is to construct a membrane with boundary Γ and mean curvature zero. In some cases, this can be achieved by as follows. Given the bounding wire frame Γ, consider some initial surface $S(t = 0)$ whose boundary is Γ. Let $S(t)$ be the family of surfaces parameterized by t obtained by allowing the initial surface $S(t = 0)$ to evolve under mean curvature, with boundary always given by Γ. Defining the surface S by $S = \lim_{t \to \infty} S(t)$, one expects that the surface S will be a minimal surface for the boundary Γ. Several computational approaches exist to construct such minimal surfaces based on this approach, including K. Brakke's Surface Evolver program [16].

A level set approach to this problem rests on embedding the motion of the surface towards its minimal energy as the zero level set of a higher dimensional function. Thus, given an initial surface $S(0)$ passing through Γ, construct a family of neighboring surfaces by viewing $S(0)$ as the zero level set of some function ϕ over all of R^3. Using the level set equation (Eqn. 24), evolve ϕ according to the speed law $F(\kappa) = -\kappa$. Then a possible minimal surface S will be given by

$$S = \lim_{t \to \infty} \{x | \phi(x, t) = 0\} \tag{87}$$

The difficult challenge with the above approach is to ensure that the evolving zero level set always remains attached to the boundary Γ. This is accomplished by Chopp [28] by creating a set of boundary conditions on those grid points closest to the wire frame that link together the neighboring values of ϕ to force the level set $\phi = 0$ through Γ. Thus, the problem turns into one of constrained level set problem; we track mean curvature flow with the constraint the evolving zero level set remained attached to the front.

In Figure 37, taken from [28], the minimal surface spanning two rings each of radius 0.5 and at positions $x = \pm.277259$ is computed. A cylinder spanning the two rings is taken as the initial level set $\phi = 0$. A $27 \times 47 \times 47$ mesh with space step 0.025 is used. The final shape is shown in from several different angles in Figure 37.

Next, in Figure 38 (again taken from [28]), this same problem is computed, but the rings are placed far enough apart so that a catenoid solution cannot exist. Starting with

| $Time = 0.0$ | $Time = 0.41$ | $Time = 0.42$ | $Time = 0.50$ |
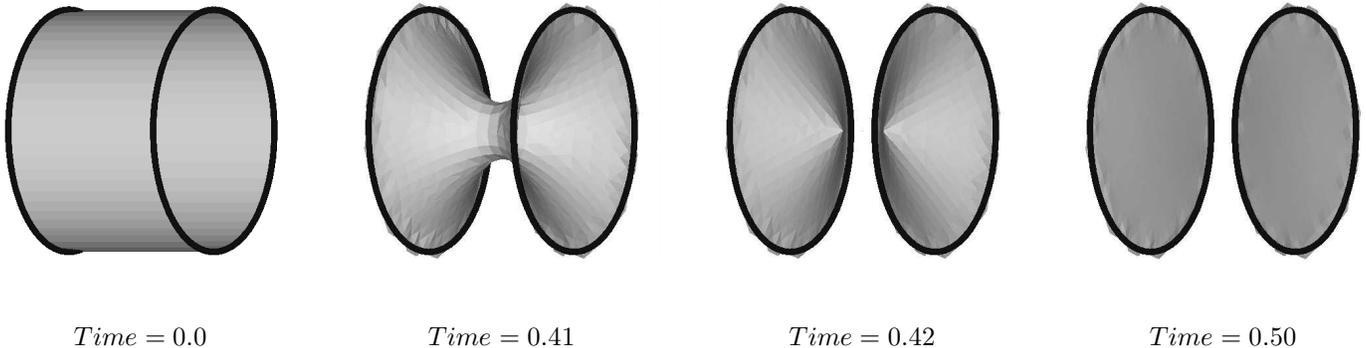
Figure 38: Splitting of Catenoid

a cylinder as the initial surface, the evolution of this surface is computed as it collapses under mean curvature while remaining attached to the two wire frames. As the surface evolves, the middle pinches off and the surface splits into two surfaces, each of which quickly collapses into a disk. The final shape of a disk spanning each ring is indeed a minimal surface for this problem. This example illustrates one of the virtues of the level set approach. No special cutting or *ad hoc* decisions are employed to decide when to break the surface. Instead, viewing the zero level set as but one member of a family of flowing surfaces allows this smooth transition. Further results may be found in [28].

21.2   Shape Detection/Recovery

Imagine that one is given an image. The goal in *shape detection/recovery* is to extract a particular shape from that image; here, "extract" means to produce a mathematical description of the shape which can be used in a variety of forms. The work on level set techniques applied to shape recovery described here was first presented in [78]; further work using the level set scheme in the context of shape recovery may be found in [79, 76, 73, 23]. We refer the interested reader to those papers for motivation, details, and a large number of examples.

Imagine that we are given an image, with the goal of isolating a shape within the image. Our approach (see [78]) is motivated by the active force contour/snake approach to shape recovery given by Kass, Witkin, and Terzopoulos [61]. Consider a speed function of the form $1 - \epsilon\kappa$ $(1 + \epsilon\kappa)$, where ϵ is a constant. As discussed earlier, the constant acts as an advection term, and is independent of the moving front's geometry. The front uniformly expands (contracts) with speed 1 (-1) depending on the sign, and is analogous to the inflation force defined in [33]. The diffusive second term ϵK depends on the geometry of the front and smooths out the high curvature regions of the front. It has the same regularizing effect on the front as the internal deformation energy term in thin-plate-membrane splines [61].

Our goal now is to define a speed function from the image data that acts as a halting criterion for this speed function. We multiply the above speed function by the term:

$$k_I(x, y) = \frac{1}{1 + |\nabla G_\sigma * I(x, y)|}, \tag{88}$$

where the expression $G_\sigma * I$ denotes the image convolved with a Gaussian smoothing filter whose characteristic width is σ. The term $|\nabla G_\sigma * I(x, y)|$ is essentially zero except

where the image gradient changes rapidly, in which case the value becomes large. Thus, the filter $k_I(x, y)$ is close to unity away from boundaries, and drops to zero near sharp changes in the image gradient, which presumably corresponds to the edge of the desired shape. In other words, the filter function anticipates steep changes in the image gradient, and stops the evolving front from passing out of the desired region.

Thus the algorithm works as follows. A small front (typically a circle) is started inside the desired region. This front then grows outwards and is stopped at the shape boundary by the filter term which drops the value of the speed function F to zero.

There are several desirable aspects of this approach.

- The initial front can consist of many fronts; due to the topological capabilities of the level set method, these fronts will merge into a single front as it grows into the particular shape.

- The front can follow intricate twists and turns in the desired boundary.

- Use of narrow band techniques makes the algorithm very fast.

- The technique can be used to extract three-dimensional shapes as well by initializing in a ball inside the desired region.

- Small isolated spots of noise where the image gradient changes substantially are ignored; the front propagates *around* these points and closes back in on itself and then disappears.

As a demonstration, level set shape recovery techniques are applied to the difficult problem of extracting the left and right ventricles of the heart. In these calculations, taken from [77], the problem is initialized by simultaneously tagging both the left and the right ventricle; the right ventricle is found by the evolving front, as is the left ventricle. Note that in the evolution of the right ventricle front, the papillary muscle is also found, see Figure 39. This feature is found by initializing with a single contour which wraps around the papillary muscle and separates into an inner ring and outer ring. After the outer wall of the left and right ventricles are recovered, the outer wall of the right ventricle is extracting; this is done by temporarily relaxing the stopping criteria, and allowing the front to move past the inner wall of the right ventricle. Once this occurs, the stopping criteria is turned back on, and the front expands until the outer wall is found.

This technique for shape detection/recovery can be performed in three dimensions if three-dimensional data is available. For details of this and related work, see [78, 79, 77].
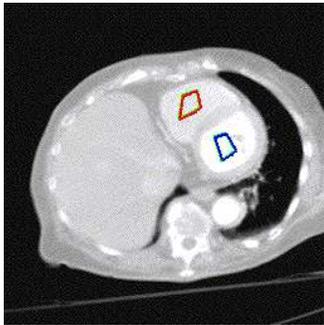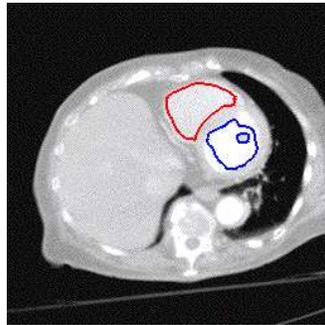
(a) Initialization   (b) Intermediate Stage   (c) End of Stage One
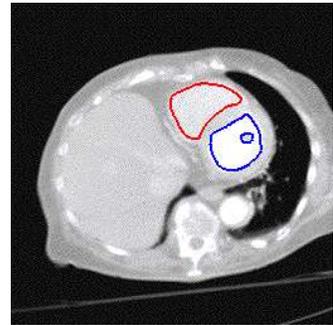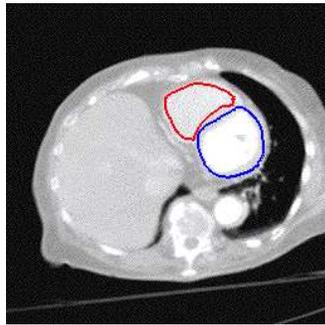
(d) Intermediate Stage Two   (e) End of Stage Two
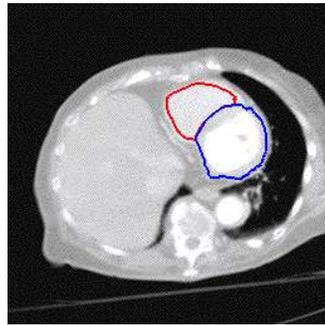
Figure 39: Shape Extraction From Heart Data

22 Applications of the Fast Marching Level Set Method

In the case of a monotonically advancing front whose speed in the normal direction depends only on position, we have previously seen that this can be converted into a stationary time problem. Furthermore, we have developed a fast marching algorithm for solving the Eikonal equation associated with this problem. Here, we show two applications of this technique.

22.1 Shape-from-shading

Suppose we illuminate a non-self-shadowing surface from a single point light source. At each point of the surface, one can define the brightness map I which depends on the reflectivity of the surface and the angle between the incoming light ray and the surface normal. Points of the surface where the normal is parallel to the incoming beam are brightest; those where the normal is almost orthogonal are darkest (again, we rule out surfaces that are self-shadowing). The goal of *shape-from-shading* is to reconstruct the surface from its brightness function I. A good starting point is the work

We point out right away that the problem as posed does not have a unique solution. For example, imagine a beam coming straight down; it is impossible to differentiate a surface from its mirror image from the brightness function. That is, a deep valley could also be a mountain peak. Other ambiguities can exist, we refer the reader to Rouy and Tourin[96] and Kimmel and Bruckstein [65]. Nonetheless, in its simplest form the shape-from-shading problem provides a simple example of an Eikonal equation that can be solved using our fast marching level set method.

We begin by considering a surface $T(x, y)$; the surface normal is then given by

$$n = \frac{(-T_x, -T_y, 1)}{(|\nabla T|^2 + 1)^{1/2}}. \tag{89}$$

Let (α, β, γ) be the direction from the light source. In the simplest case of a Lambertian surface, the brightness map is given in a very simple form by

$$I(x, y) = (\alpha, \beta, \gamma) \cdot n. \tag{90}$$

Thus, the shape-from-shading problem is to reconstruct the surface $T(x, y)$ given the brightness map $I(x, y)$.

Consider the simplest case, namely that in which the light comes from straight down. Then the light source vector is $(0, 0, 1)$, and we have

$$I(x, y) = \frac{1.}{(|\nabla T|^2 + 1)^{1/2}}. \tag{91}$$

Rearranging terms, we then have an Eikonal equation for the surface, namely

$$|\nabla T| = \sqrt{\frac{1}{I^2} - 1}. \tag{92}$$

where n is the normal to the surface. We still need initial conditions for this problem Let us imagine that at extrema of T we know the values of T. Then we can construct a viable solution surface using our fast marching method.

To demonstrate, we start with a given surface, first compute the brightness map I, and then reconstruct the surface by solving the above Eikonal equation. In Figure 40, we show a paraboloid surface of the form $T = 3. - 3(x^2 + y^2)$. In Figure 40(a) we show

(a) Original   (b) Brightness   (c) Reconstructed
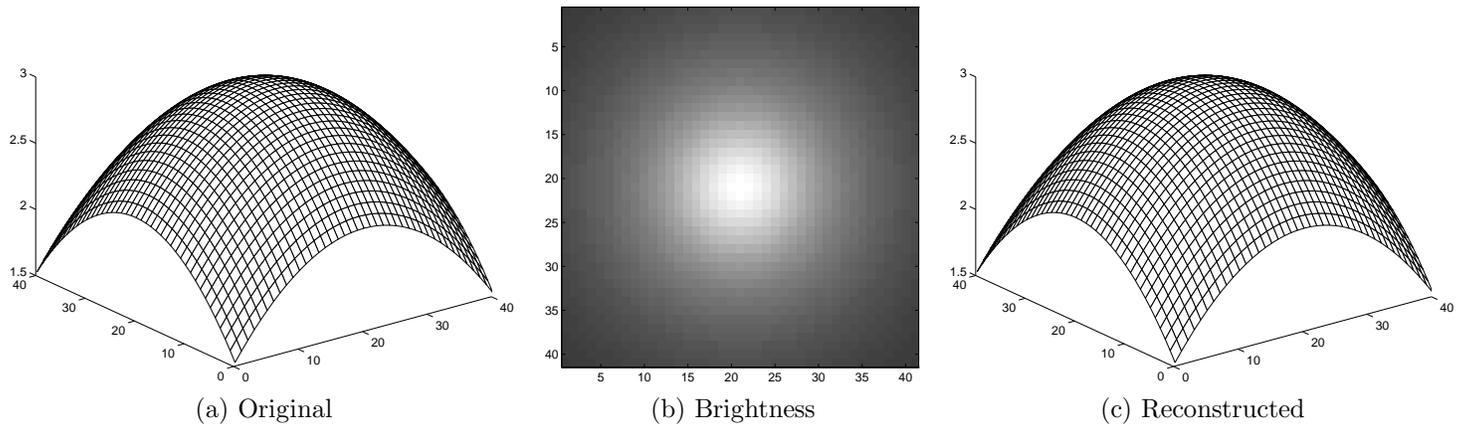
Figure 40: Shape-from-shading reconstruction of Paraboloid surface

(a) Original   (b) Brightness   (c) Reconstructed
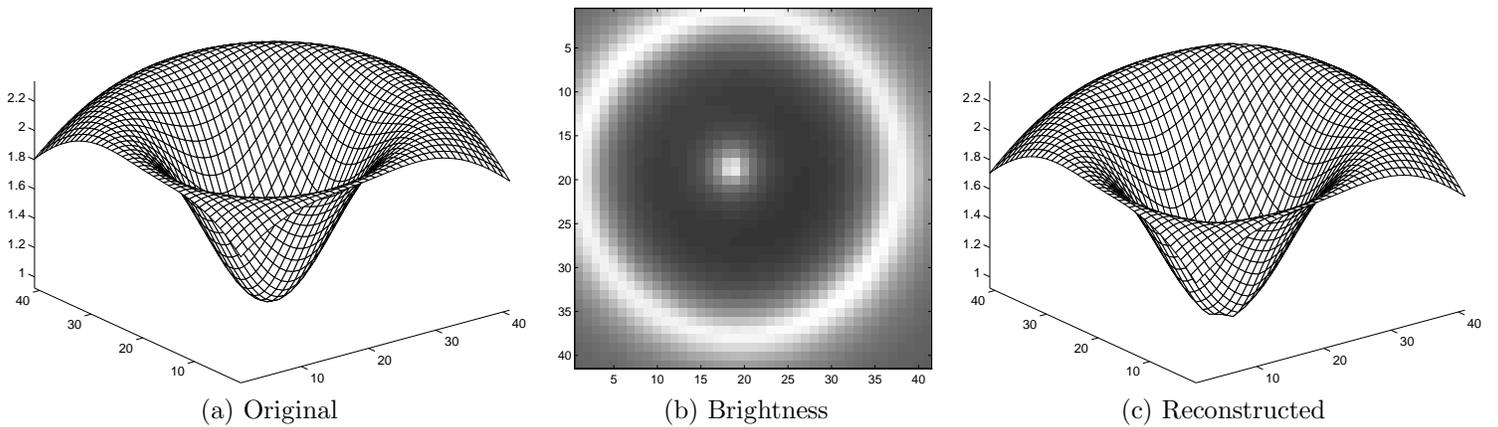
Figure 41: Shape-from-shading reconstruction of Double Gaussian surface

the original surface, in Figure 40()b we show the brightness map $I(x, y)$, and in Figure 40(c) we show the reconstructed surface. This surface is "built" by setting $T = 3$ at the point where the maximum is obtained, and then solving the Eikonal equation.

As a more complex example, we use a double Gaussian function of the form

$$T(x, y) = 3e^{-(x^2+y^2)} - 2e^{-20((x-.05)^2+(y-.05)^2)} \qquad (93)$$

Once again, we compute the brightness map and then reconstruct the surface, see Figure 41.

We have barely touched the topic of shape-from-shading; in the case of multiple extrema and non-vertical light sources, more care must be taken, and we refer the interested reader to the above sources. Nonetheless, the fast marching level set algorithm is extremely effective for these problems.

22.2 Photolithography Development

One component process in the manufacturing of microchips is the stage of *lithography development*, see Chapter 17. In this process, the resist properties of a material have been altered due to exposure to a beam which has been partially blocked by a pattern mask.
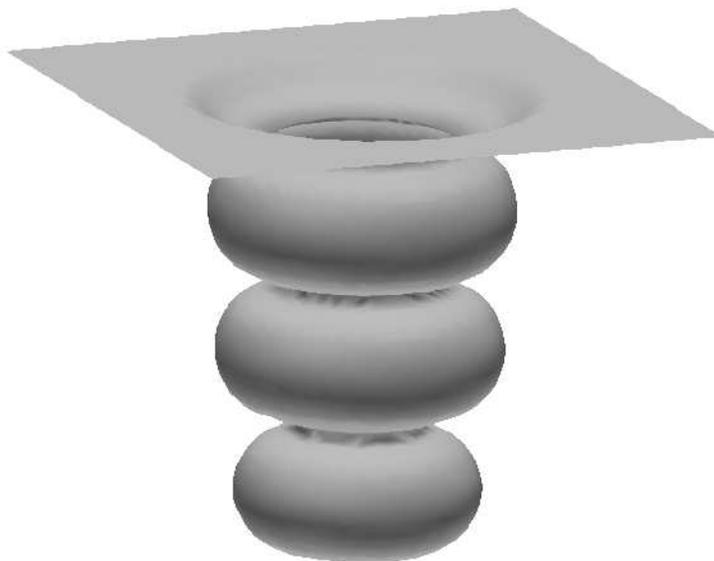
Figure 42: Lithographic development on 50x50x50 grid

Grid Size	Time to Load Rate File	Time to Propagate Front	Total Time
50x50x50	0.1 secs	0.7 secs	0.8 secs
100x100x100	1.2 secs	8.2 secs	9.4 secs
150x150x150	3.9 secs	37.8 secs	41.7 secs
200x200x200	9.0 secs	80.0 secs	89 secs

Figure 43: Timings for Development to T=10: Sparc 10

The material is then "developed", which means the material with less resistivity is etched away. While the process is discussed in more detail in the next section, at this point we simply note that the problem reduces to that of following an initially plane interface propagating downwards in three dimensions, where the speed in the normal direction is given as a supplied rate function at each point. The speed $F = F(x, y, z)$ depends only on position, however it may change extremely rapidly. The goal in lithography development is to track this evolving front. In order to develop realistic structures in three-dimensional development profiles, a grid of size 300x300x100 is not unreasonable; hence a fast algorithm is required to perform the development step.

Start with a flat profile at height $z = 1$ in the unit cube centered at $(.5, .5, .5)$ and follow the evolution of the interface downwards with speed given by the model Gaussian rate function

$$F(x, y, z) = e^{-64(r^2)}(\cos^2(12z) + .01) \tag{94}$$

where $r = \sqrt{(x - .5)^2 + (y - .5)^2}$. This rate function F corresponds to effect of standing waves which change the resist properties of the material, and causes sharp undulations and turns in the evolving profile. In Figure 42, we show the profile etched out by such an initial state; the calculation is carried out until $T = 10$.

In Figure 43, we give timings for a parameter study on a Sparc10 for the speed function $F = e^{-64(r^2)}(\cos^2(6z) + .01)$. We note that loading the file containing the model Gaussian rate function F is a significant proportion of the total compute time.

Further details of the application of our fast marching level set method may be found

in [111, 110, 116]

23 A Final Example: Etching and Deposition in the Microfabrication of Semiconductor Devices

23.1 Background

A goal of numerical simulations in microfabrication of semi-conductor devices is to model the process by which silicon devices are manufactured. Here, we briefly summarize the stages involved. First, a single crystal ingot of silicon is extracted from molten pure silicon. This silicon ingot is then sliced into a several hundred thin wafers, each of which is then polished to a smooth finish. A thin layer of crystalline is then oxidized, a light sensitive "photoresist" that is sensitive to light is applied, and then the wafer is covered with a pattern mask which shields part of the photoresist. This pattern mask contains the layout of the circuit itself. Under exposure to a light or an electron beam, the exposed photoresist polymerizes and hardens, leaving an unexposed material which is then etched away in dry etch process, revealing bare silicon dioxide layer. Ionized impurity atoms such as boron, phosphorus and argon are then implanted into the pattern of the exposed silicon water, and silicon dioxide is deposited at reduced pressure in a plasma discharge from gas mixtures at a low temperature. Finally, thin films such as aluminum are deposited by processes such as plasma sputtering, and contacts to the electrical components and component interconnections are established. The result is a device that carries the desired electrical properties.

The above processes produce considerable change in the surface profile as it undergoes the stages of etching, deposition, and photolithography. This problem is known as the "surface topography problem" in microfabrication, and is controlled by a large collection by physical effects, including the visibility of the etching/deposition source at each point of the evolving profile, surface diffusion along the front, non-convex sputter laws which produce faceting, shocks and rarefactions, material-dependent discontinuous etch rates, and masking profiles.

The underlying physical effects involved in etching, deposition and lithography are quite complex; excellent overviews are due to Neureuther and his group, see [56, 101, 102, 125, 126], as well as [20, 21, 22, 80, 94]. The effects may be summarized briefly as follows:

- *Deposition:* Particles are deposited on the surface, which causes build-up in the profile. The particles may either isotropically condense from the surroundings (known as chemical or "wet" deposition), or be deposited from a source. In the latter case, we envision particles leaving the source and depositing on the surface; the main advantage of this approach is increased control over the directionality of surface deposition. The rate of deposition, and hence growth of the layer, may depend on source masking, visibility effects between the source and surface point, angle-dependent flux distribution of source particles, the angle of incidence of the particles relative to the surface normal direction, reflection of deposited particles, and surface diffusion effects.

- *Etching:* Particles remove material from the evolving profile boundary. The material may be isotropically removed, known again as chemical or "wet" etching, or chipped away through reactive ion etching, also known as "ion milling". Similar to deposition, the main advantage of reactive ion etching is enhanced directionality, which becomes increasingly important as device sizes decrease substantially and

etching must proceed in vertical directions without affecting adjacent features. The total etch rate consists of an ion-assisted rate and a purely chemical etch rate due to etching by neutral radicals, which may still have a directional component. As in the above, the total etch rate due to wet and directional milling effects can depend on source masking, visibility effects between the source and surface point, angle-dependent flux distribution of source particles, the angle of incidence of the particles relative to the surface normal direction, reflection/re-emission of etching/million particles, and surface diffusion effects.

- *Lithography:* As discussed earlier, the underlying material is created by an electromagnetic wave which alters the resist property of the material. The aerial image is found, which then determines the amount of crosslinking at each point in the material which then produces the etch/resist rate at each point of the material. A profile is then etched into the material, where the speed of the profile in its normal direction at any point is given by the underlying etch rate. The key factors that determine the evolving shape are the etch/resist profile and masking effects.

In the final analysis, the above reduces to our familiar problem of tracking the boundary of moving interface moving under a speed function F. Abstractly, we may write

$$F = F_{Deposition/Etching} + F_{Lithography}. \tag{95}$$

Of course, all effects do not take place at once; however the design of the numerical algorithm allows various combinations of terms to be "turned on" during any time step of the surface advancement. For details and additional calculations of level set methods applied to microfabrication, see [2, 3, 4].

23.2 Results

23.2.1 Etching/Deposition

We begin in Figure 44 with a deposition source above a trench, where deposition material emitted from a line source from the solid line above the trench. In this experiment, the deposition rate is the same in all directions. The effects of shadowing are considered. Figure 44a shows results for 40 computational cells along the width the compute region (between the two vertical dashed lines); Figure 44b has 80 cells, and Figure 44c has 160 cells. The time step for all three calculations is $\Delta t = .00625$. The calculations are performed with a narrow band tube width of 6 cells on either side of the front. There is little change between the calculation with 80 cells and the one with 160 cells, indicating that the solution is converged. As the walls pinch toward each other, the seen visible angle decreases and the speed diminishes.

23.2.2 Ion-Milling: Non-convex Sputter Laws

A more sophisticated set of examples arises in simulations (for example, of ion milling) in which the normal speed of the profile depends the angle of incidence between the surface normal and the incoming beam. This yield function is often empirically fit from experiment, and has been observed to cause such effects as faceting at corners, see [69, 62]. As shown in [2, 3], such yield functions can often give rise to non-convex Hamiltonians, in which case alternative schemes must be used. To study this phenomenon, in Figure 45 we consider several front motions and their effects on corners. We envision an etching beam coming down in the vertical direction. In the cases under study here, the angle θ refers to the angle between the surface normal and the positive vertical. For this set of calculations, in order to focus on the geometry of sputter effects on shocks/rarefaction fan
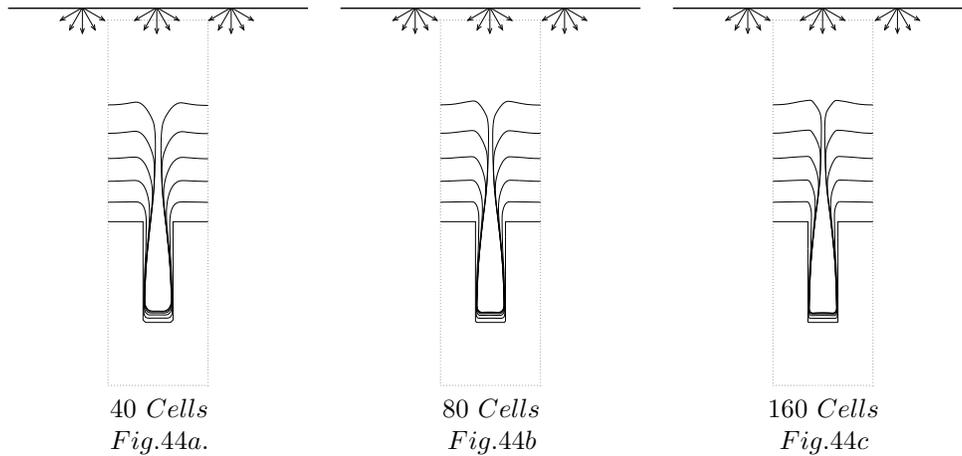
Figure 44: Source Deposition Into Trench

development, visibility effects are ignored. The calculations are made using the schemes for non-convex Hamiltonians described earlier. Following our usual notation, let $F(\theta)$ be the speed of the front in direction normal to the surface.

In column A, the effects of purely isotropic motion are shown, thus the yield function is $F = 1$. Located above the yield graph are the motions of a downwards square wave. In column B, the effects of directional motion are shown, thus the yield function is $F = \cos(\theta)$. In this case, the horizontal components on the profile do not move, and vertical components move with unit speed. In column C, the effects of a yield function of the form $F = [1 + 4\sin^2(\theta)]\cos(\theta)$ are shown.

The results of these calculations are given in Figure 45. The results show that the effects of angle dependent yield functions are pronounced. In Column A the isotropic rate produces smooth corners, correctly building the necessary rarefaction fans in outward corners and entropy satisfying shocks in inward corners as discussed and analyzed in [105, 106]. In Column B, the directional rate causes the front to be essentially translated upwards, with minimal rounding of the corners. In Column C, the yield function results in faceting of inward corners where shocks form and sharp corners in the construction of rarefaction fans.

23.2.3 Discontinuous Etch Rates

Next, we study the effects of etching through different materials. In this example, the etch rates are discontinuous, and hence sharp corners develop in the propagating profile. The results of these calculations are show in Figure 46. A top material masks a lower material, and the profile etches through the lower material first and underneath the upper material. The profile depends on the ratio of the etch rates. In Figure 46a, the two materials have the same etch rate, and hence the front simply propagates in its normal direction with unit speed, regardless of the which material it is passing through. In Figure 46b, the bottom material etches four times faster that the top; in Figure 46c, the ratio is ten to one, Finally, in Figure 46d, the ratio is forty to one, in which case the top material almost acts like a mask.

$$F(\theta) = 1 \qquad F = \cos(\theta) \qquad F = [1 + 4\sin^2(\theta)]\cos(\theta)$$
$$Column\ A \qquad\qquad Column\ B \qquad\qquad Column\ C$$
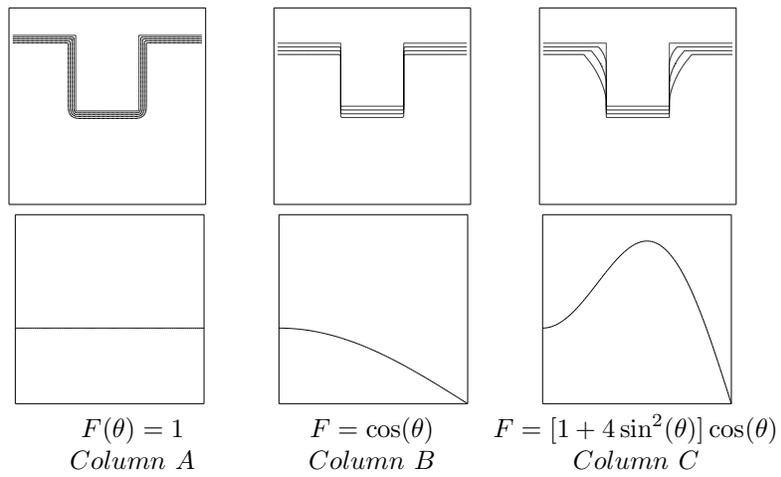
Figure 45: Effect of Different Yield Functions: Non-convex Scheme

$$Ratio\ = 1:1 \qquad Ratio\ = 4:1 \qquad Ratio\ = 10:1 \qquad Ratio\ = 40:1$$
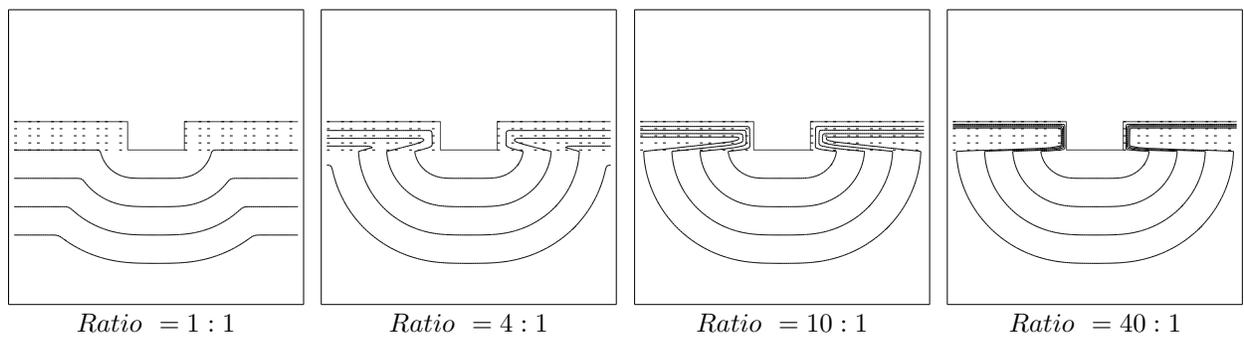
Figure 46: Etch Ratio = Bottom Material Rate to Top Material Rate

23.2.4   Simultaneous Etching and Deposition

Next, a parameter study of simultaneous etching and deposition is taken from [4]. The speed function is

$$F = (1 - \alpha)F_{etch} + \alpha F_{Deposition} \tag{96}$$

where

$$F_{etch} = (5.2249\cos\theta - 5.5914\cos^2\theta + 1.3665\cos^4\theta) \quad F_{Deposition} = \beta F_{Isotropic} + (1-\beta)F_{Source} \tag{97}$$

Visibility effects are considered in all terms except isotropic deposition. Figure 47 shows the results of varying α and β between 0 and 1.

23.3   Three-Dimensional Simulations

Finally, a three-dimensional example of a non-convex sputter yield law is applied to an indented saddle, which gives rise to faceting as shown in Figure 48. Complete details of the above and a large variety of simulations of etching, deposition, and lithography development may be found in [2, 3, 4].

24   Other Work

The range of level set techniques extends far beyond the work covered here. Here, we point the reader to some additional topics.

On the theoretical side, considerable analysis of level set methods has been performed in recent years, see, for example, [15, 38, 41, 42, 43, 44, 26, 50, 51, 6]. This work has concentrated on many aspects, including questions of existence and uniqueness, pathological cases, extensions of these ideas to fronts of co-dimension greater than one (such as evolving curves in three dimensions), coupling with diffusion equations, links between the level set technique and Brakke's groundbreaking original varifold approach.

On the theoretical/numerical analysis side, level set techniques exploit the considerable technology developed in the area of viscous solutions to Hamilton-Jacobi equations, see the work in [9, 35, 36, 37, 71].

A wide range of applications relate to level set methods, including work on minimal arrival times by Falcone [45], flame propagation work by Zhu and Ronney [129], a wide collection of applications from computer vision by Kimmel [64], gradient flows applied to geometric active contour models [63], work on affine invariant scale space [99], and some related work on the scalar wave equation [47]. We also refer the reader to the collection of papers from the International Conference on Mean Curvature Flow [18].

References

[1] Adalsteinsson, D., and Sethian, J.A., *A Fast Level Set Method for Propagating Interfaces*, J. Comp. Phys., Vol. 118(2), pp. 269–277, May 1995.

[2] Adalsteinsson, D., and Sethian, J.A., *A Unified Level Set Approach to Etching, Deposition and Lithography I: Algorithms and Two-dimensional Simulations*, J. Comp. Phys. 120, 1, pp. 128-144, 1995.
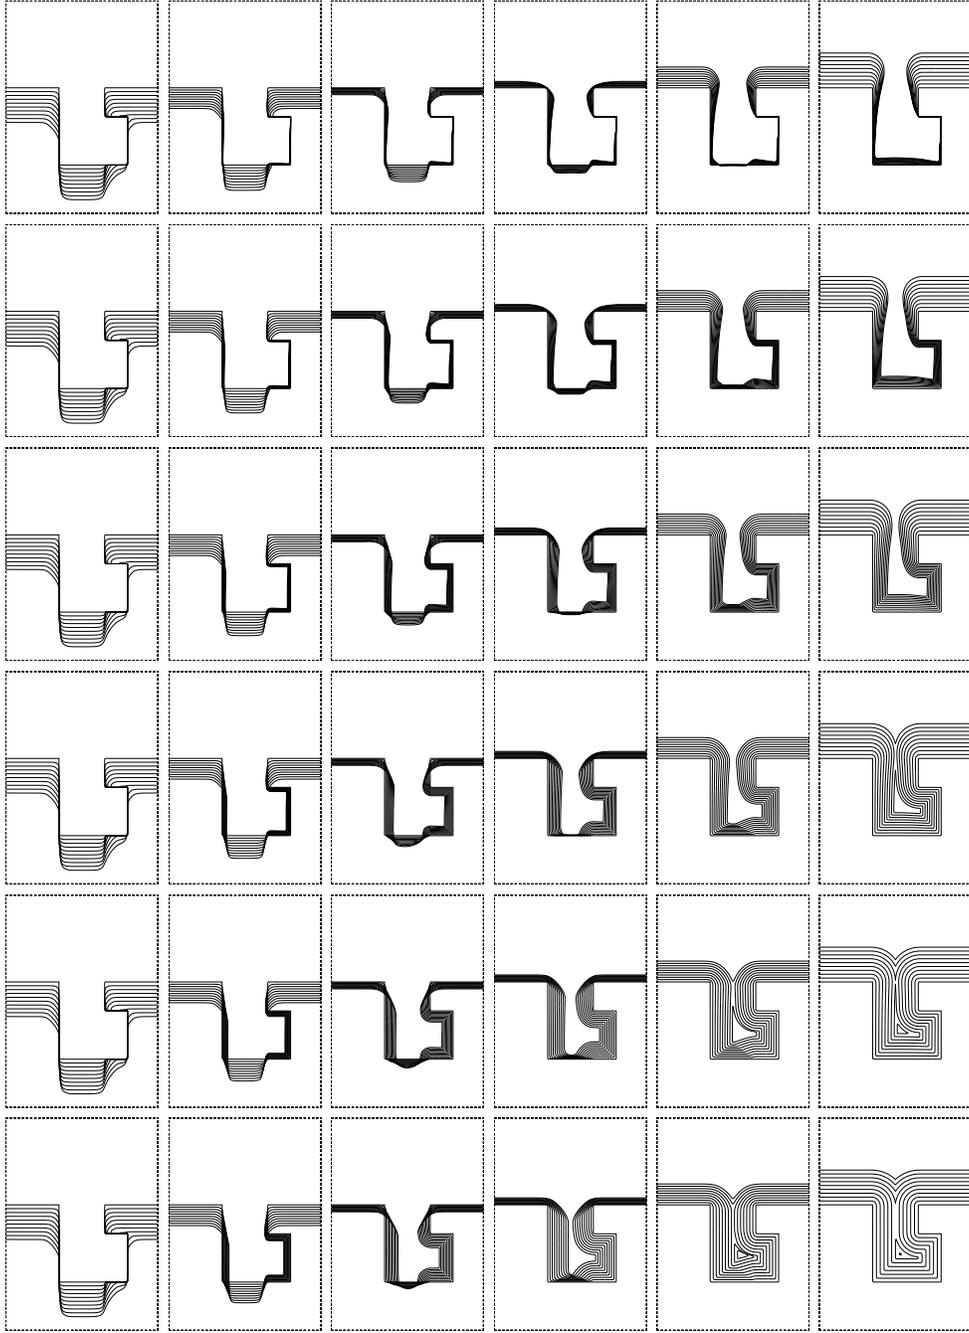
$$F = (1 - \alpha)F_{etch} + \alpha F_{Deposition}$$

$$F_{etch} = (5.2249\cos\theta - 5.5914\cos^2\theta + 1.3665\cos^4\theta)\cos\theta$$

$$F_{Deposition} = \beta F_{Isotropic} + (1 - \beta)F_{Source}$$
$$\alpha\ Increases\ From\ Left\ to\ Right$$
$$\beta\ Increases\ From\ Bottom\ to\ Top$$

Figure 47: Simultaneous Etching and Deposition

$Initial \quad Shape : T = 0$ $F = [1 + 4\sin^2(\theta)]\cos(\theta) \quad T = 8$ $Final \quad Rotated$
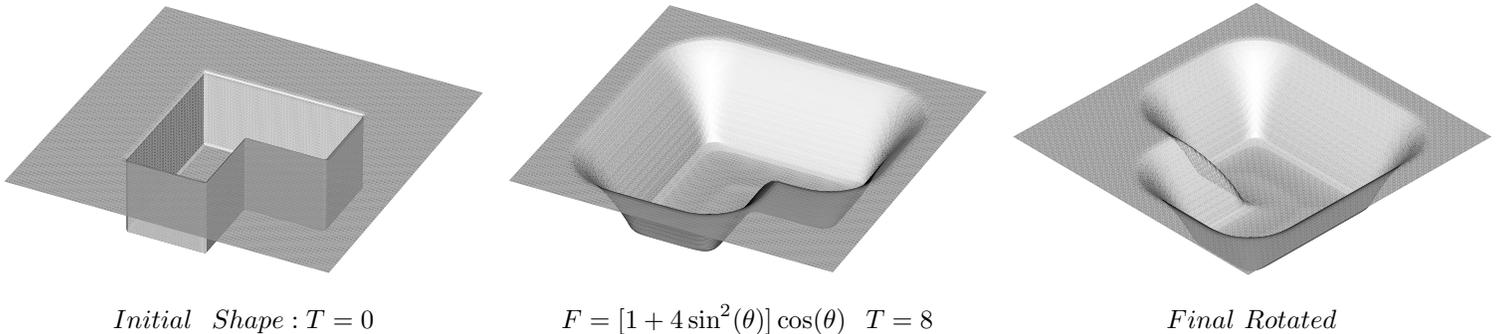
Figure 48: Downward Saddle Under Sputter Etch

[3] Adalsteinsson, D., and Sethian, J.A., *A Unified Level Set Approach to Etching, Deposition and lithography II: Three-dimensional Simulations*, J. Comp. Phys., Vol. 122, No. 2, pp.348-366, 1995.

[4] Adalsteinsson, D., and Sethian, J.A., *A Unified Level Set Approach to Etching, Deposition and Lithography III: Complex Simulations and Multiple Effects*, to be submitted, J. Comp. Phys., 1995.

[5] Alvarez, L., Lions, P.L., and Morel, M., *Image Selective Smoothing and Edge detection by Nonlinear Diffusion. II,* SIAM Journal on Numerical Analysis, Vol. 29(3), pp. 845–866, 1992.

[6] Ambrosio, L., and Sonar, H.M., *Level Set Approach to Mean Curvature Flow in Arbitrary Codimension*, preprint, 1994.

[7] Angenent, S.,*Shrinking Doughnuts*, in "Proceedings of Nonlinear Diffusion Equations and Their Equilibrium States, 3, (Edited by N.G. Lloyd et. al.), Birkhauser, Boston, 1992.

[8] Barles, G., *Remarks on a Flame Propagation Model*, INRIA Report 464, 1985.

[9] Barles, G., *Discontinuous Viscosity Solutions of First Order Hamilton-Jacobi Equations: A Guided Visit*, Non-linear Analysis: Theory, Methods, and Applications, 20, 9, pp. 1123-1134, 1993.

[10] Barles, G., and Souganidis, P.E., *Convergence of Approximation Schemes for Fully Non-linear Second Order Equations*, Asymptotic Anal., 4, pp.271-283, 1991.

[11] Bardi, M., and Falcone, M., *An Approximation Scheme for the Minimum Time Function*, SIAM J. Control Optim, 28, pp. 950-965, (1990),

[12] Bell, J.B., Colella, P., and Glaz, H.M., *A Second-Order Projection Method for the Incompressible Navier-Stokes Equations*, J. Comp. Phys., 85, pp. 257-283, 1989.

[13] Berger, M. and Colella, P., *Local Adaptive Mesh Refinement for Shock Hydrodynamics*, J. Comp., Phys., 1, 82, pp. 62-84, 1989.

[14] Brackbill, J.U., Kothe, D.B., and Zemach, C., *A Continuum Method for Modeling Surface Tension*, J. Comp. Phys., 100, pp. 335-353, 1992.

[15] Brakke, K.A., *The Motion of a Surface by Its Mean Curvature*, Princeton University Press, Princeton, New Jersey, 1978.

[16] Brakke , K.A., Surface Evolver Program, Research Report GCC 17, the Geometry Supercomputer Project, University of Minnesota, 1200 Washington Ave., South, Minneapolis, Minn., 55455, 1990.

[17] Bronsard, L., and Wetton, B., *A Numerical Method for Tracking Curve Networks Moving with Curvature Motion*, 120, 1, pp. 66087, 1995.

[18] Buttazzo, G., and Visitin, A., *Motion by Mean Curvature and Related Topics*, Proceedings of the International Conference at Trento, 1992, Walter de Gruyter, New York, 1994.

[19] Cahn, J.E., and Hilliard, J.E., Jour. Chem. Phys. 28, pp. 258, 1958.

[20] Cale, T.S., and Raupp, G.B., *Free Molecular Transport and Deposition in Cylindrical Features*, J. Vac. Sci. Tech., B, 8,4, pp. 649-655, 1990.

[21] Cale, T.S., and Raupp, G.B., *Free Molecular Transport and Deposition in Long Rectangular Trenches*, J. Appl. Phys., 68, 7, pp.3645-8652, 1990.

[22] Cale, T.S., and Raupp, G.B., *A Unified Line-of-Sight Model of Deposition in Rectangular Trenches*, J. Vac. Sci. Tech., B, 8,6, pp. 1242-1248, 1990.

[23] Caselles, V., Catte, F., Coll, T., and Dibos, F., *A Geometric Model for Active Contours in Image Processing*, Internal report no. 9210, CEREMADE, Université de Paris-Dauphine, France.

[24] Castillo, J.E., *Mathematical Aspects of Grid Generation.* Frontiers in Applied Mathematics 8, SIAM Publications, 1991.

[25] Chang, Y.C., Hou, T.Y., Merriman, B., and Osher, S.J., *A Level Set Formulation of Eulerian Interface Capturing Methods for Incompressible Fluid Flows*, submitted for publication, Jour. Comp. Phys., 1994.

[26] Chen, Y., Giga, Y., and Goto, S., *Uniqueness and Existence of Viscosity Solutions of Generalized Mean Curvature Flow Equations*, J. Diff. Geom, Vol. 33, 749, 1991.

[27] Chopp, D.L., *Numerical Computation of Self-Similar Solutions for Mean Curvature Flow* J. Exper. Math., Vol. 3, No. 1, pp.1-15, 1994.

[28] Chopp, D.L., *Computing Minimal Surfaces via Level Set Curvature Flow*, Jour. of Comp. Phys., Vol. 106, pp. 77–91, 1993.

[29] Chopp, D.L., and Sethian, J.A,. *Flow Under Curvature: Singularity Formation, Minimal Surfaces, and Geodesics*, Jour. Exper. Math., Vol. 2(4), pp. 235–255, 1993.

[30] Chorin, A.J., *Flame Advection and Propagation Algorithms*, J. Comp. Phys., Vol. 35, pp. 1-11, 1980.

[31] Chorin, A.J., *Numerical Solution of the Navier-Stokes Equations*, Math. Comp., 22, pp. 745, 1968.

[32] Chorin, A.J., *Numerical Study of Slightly Viscous Flow*, J. Fluid Mech., 57, pp. 785-796, 1973.

[33] Cohen, L.D., *On Active Contour Models and Balloons*, Computer Vision, Graphics, and Image Processing, Vol. 53, No. 2, pp. 211–218, March 1991.

[34] Colella, P., and Puckett, E.G., *Modern Numerical Methods for Fluid Flow*, Lecture Notes, Department of Mechanical Engineering, University of California, Berkeley, California, 1994.

[35] Crandall, M.G., Evans, L.C., and Lions, P-L., *Some Properties of Viscosity Solutions of Hamilton-Jacobi Equations*, Tran. AMS, 282, pp. 487-502, 1984.

[36] Crandall, M.G., Ishii, H., and Lions, P-L., *User's Guide to Viscosity Solutions of Second Order Partial Differential Equations*, Bull. AMS, 27/1, pp.1-67, 1992.

[37] Crandall, M.G., and Lions, P-L., *Viscosity Solutions of Hamilton-Jacobi Equations*, Tran. AMS, 277, pp. 1-43, 1983.

[38] Ecker, K., Huisman, G., *Interior Estimates for Hypersurfaces Moving by Mean Curvature*, Inventiones Mathematica, 105, 3, pp. 547-569, 1991.

[39] Engquist, B., and Osher, S.J., *Stable and Entropy-Satisfying Approximations for Transonic Flow Calculations*, Math. Comp., Vol. 34, 45, 1980.

[40] Evans, L.C., Sonar, H.M., and Souganidis, P.E., *Phase Transitions and Generalized Motion by Mean Curvature*, Communications on Pure and Applied Mathematics, Vol. 45, 1097, 1992.

[41] Evans, L.C., and Spruck, J., *Motion of Level sets by Mean Curvature I*, J. Diff. Geom, Vol. 33, 635, 1991.

[42] Evans, L.C., and Spruck, J., *Motion of Level sets by Mean Curvature II*, Transactions of the American Mathematical Society, Vol. 330, 91, 1992.

[43] Evans, L.C., and Spruck, J., *Motion of Level sets by Mean Curvature III*, J. Geom. Anal. 2, pp. 121-150, 1992.

[44] Evans, L.C., and Spruck, J., *Motion of Level sets by Mean Curvature IV*, J. Geom. Anal.,5, 1, pp. 77-114, 1995.

[45] Falcone, M., *The Minimum Time Problem and Its Applications to Front Propagation*, in "Motion by Mean Curvature and Related Topics", Proceedings of the International Conference at Trento, 1992, Walter de Gruyter, New York, 1994.

[46] Falcone, M., Giorgi, T., and Loretti, P., *Level Sets of Viscosity Solutions: Some Applications to Fronts and Rendez-Vous Problems*, SIAM J. Appl. Math., 54, 5, pp. 1335-1354, 1994.

[47] Fatemi, E., Engquist, B., and Osher, S.J., *Numerical Solution of the High Frequency Asymptotic Wave Equation for the Scalar Wave Equation*, J. Comp. Phys,. 120, pp. 145-155, 1995.

[48] Gage, M., *Curve Shortening Makes Convex Curves Circular*, Inventiones Mathematica, Vol. 76, 357, 1984.

[49] Gage, M., and Hamilton, R., *The Equation Shrinking Convex Planes Curves*, J. Diff. Geom, Vol. 23, 69, 1986.

[50] Giga, Y., and Goto, S., *Motion of Hypersurfaces and Geometric Equations*, Journal of the Mathematical Society of Japan, Vol. 44, 99, 1992.

[51] Giga, Y., Goto, S., Ishii, H., *Global Existence of Weak Solutions for Interface Equations Coupled with Diffusion Equations*, SIAM J. Math. Anal., 23, N4, pp. 821-835, 1992.

[52] Grayson, M., *The Heat Equation Shrinks Embedded Plane Curves to Round Points*, J. Diff. Geom., Vol. 26, 285 (1987).

[53] Grayson, M., *A Short Note on the Evolution of Surfaces Via Mean Curvatures*, J. Diff. Geom., Vol. 58, 555 (1989).

[54] Greengard, L., and Strain, J., *A Fast Algorithm for Evaluating Heat Potentials*, Comm. Pure Appl. Math., XLIII, pp. 949-963, (1990).

[55] Harten, A., Engquist, B., Osher, S., and Chakravarthy, S.R., *Uniformly high order accurate essentially non-oscillatory schemes. III*, J. Comp. Phys., 71, 2, pp. 231-303, 1987

[56] Helmsen, J.J., *A Comparison of Three-Dimensional Photolithography Development Methods* Ph.D. Dissertation, EECS, University of California, Berkeley, 1994.

[57] Hirt, C.W., and Nicholls, B.D., *Volume of Fluid)COF) Method for Dynamics of Free Boundaries*, J. Comp. Phys., 39, pp. 201-225, 1981.

[58] Huisken, G., *Flow by Mean Curvature of Convex Surfaces into Spheres*, J. Diff. Geom., Vol. 20, 237, (1984).

[59] Ilmanen, T., *Elliptic Regularization and Partial Regularity for Motion by Mean Curvature*, Memoirs of the American Mathematical Society, 108, 520, 1994.

[60] Ilmanen, T., *Generalized Flow of Sets by Mean Curvature on a Manifold*, Indiana University Mathematics Journal, 41, 3, pp.671-705, 1992.

[61] Kass, M., Witkin, A., and Terzopoulos, D., *Snakes: Active Contour Models*, International Journal of Computer Vision, pp. 321–331, 1988.

[62] Katardjiev, I.V., Carter, G., Nobes, M.J., *Precision Modeling of the Mask-Substrate Evolution During Ion Etching*, J. Vac. Science Technology, A 6(4), pp. 2443-2450, 1988.

[63] Kichenassamy, S., Kumar, A., Olver, P., Tannenbaum, A., and Yezzi, A., *Gradient Flows and Geometric Active Contours*, 1994, ICCV, 1995.

[64] Kimmel, R., *Curve Evolution on Surfaces*, Ph.D. Thesis, Dept. of Electrical Engineering, Technion, Israel, 1995.

[65] Kimmel, R., and Bruckstein, A., *Shape from Shading via Level Sets*, Center for Intelligent Systems Report No.9209, Technion- Israel Institute of Technology, June 1992,

[66] Kimmel, R., and Bruckstein, A., *Shape Offsets via Level Sets*, Computer-Aided Design, 25, 3, pp. 154-161, 1993.

[67] Knupp, P.M., and Steinberg, S., *The Fundamentals of Grid Generation*, preprint 1993.

[68] Lax, P.D., *Hyperbolic Systems of Conservation Laws and the Mathematical Theory of Shock Waves*, SIAM Reg. Conf. Series, Lectures in Applied Math, 11, pp. 1-47, 1970.

[69] Leon, F.A., Tazawa, S., Saito, K., Yoshi, A., and Scharfetter, D.L., *Numerical Algorithms for Precise Calculation of Surface Movement in 3-D Topography Simulation*, 1993 International Workshop on VLSI Process and Device Modeling (1993 VPAD).

[70] LeVeque, R.J., *Numerical Methods for Conservation Laws*, Birkhauser, Basel, 1992.

[71] Lions, P.L. *Generalized Solution of Hamilton-Jacobi Equations*, Pittman, London,1982.

[72] Majda, A., and Sethian, J.A., *Derivation and Numerical Solution of the Equations of Low Mach Number Combustion*, Combustion Science and Technology, 42, pp. 185-205, 1984.

[73] Malladi, R., Adalsteinsson, D., and Sethian, J.A., *A Fast Level Set Algorithm for 3D Shape Recovery*, submitted, IEEE Transactions on PAMI,. 1995.

[74] Malladi, R., and Sethian, J.A., *Image Processing: Flows under Min/Max Curvature and Mean Curvature*, in press, Graphical Models and Image Processing, July 1995.

[75] Malladi, R., and Sethian, J.A., *Image Processing via Level Set Curvature Flow*, Proc. Natl. Acad. of Sci., USA, Vol. 92(15), pp. 7046–7050, July 1995.

[76] Malladi, R., and Sethian, J.A., *A Unified Approach for Shape Segmentation, Representation, and Recognition*, Report LBL-36069, Lawrence Berkeley Laboratory, University of California, Berkeley, August 1994.

[77] Malladi, R., and Sethian, J.A., *A Unified Approach to Noise Removal, Image Enhancement, and Shape Recovery*, in press, IEEE Image Processing, July 1995.

[78] Malladi, R., Sethian, J.A., and Vemuri, B.C., *Evolutionary Fronts for Topology-independent shape Modeling and Recovery*, in Proceedings of Third European Conference on Computer Vision, LNCS Vol. 800, pp. 3–13, Stockholm, Sweden, May 1994.

[79] Malladi, R., Sethian, J.A., and Vemuri, B.C., *Shape Modeling with Front Propagation: A Level Set Approach*, IEEE Trans. on Pattern Analysis and Machine Intelligence, Vol. 17(2), pp. 158–175, Feb. 1995.

[80] McVittie, J.P.; Rey, J.C.; Bariya, A.J., and others *SPEEDIE: A Profile Simulator for Etching and Deposition*, Proceedings of the SPIE - The International Society for Optical Engineering, 1991, vol.1392:126-38.

[81] Merriman, B., Bence, J., and Osher, S.J., *Motion of Multiple Junctions: A Level Set Approach*, to Jour. Comp. Phys., Vol. 112, No. 2, pp. 334-363, 1994.

[82] Milne, B., and Sethian, J.A., *Adaptive Mesh Refinement for Level Set Methods for Propagating Interfaces*, to be submitted for publication, J. Comp. Phys., 1995,

[83] Mulder, W., Osher, S.J., Sethian, J.A., *Computing Interface Motion in Compressible Gas Dynamics*, Jour. Comp. Phys., Vol. 100, pp. 209-228, 1992.

[84] Mullins, W. W., and Sekerka, R. F., *Morphological Stability of a Particle Growing by Diffusion or Heat Flow*, Jour. Appl. Phys., 34, pp. 323-329, 1963

[85] Noh, W. and Woodward, P.: *A Simple Line Interface Calculation.* Proceedings, Fifth International Conference on Fluid Dynamics, A.I. vn de Vooran and P.J. Zandberger, Eds.: Springer-Verlag, 1976.

[86] Osher, S., and Rudin, L.I., *Feature-oriented Image Enhancement Using Shock Filters*, SIAM J. Num. Anal., Vol. 27, pp. 919–940, 1990.

[87] Osher, S., and Rudin, L.I., *Rapid Convergence of Approximate Solutions of Shape-From-Shading*, to appear, 1992.

[88] Osher, S., and Sethian, J.A., *Fronts Propagating with Curvature Dependent speed: Algorithms Based on Hamilton-Jacobi Formulation*, Journal of Computational Physics, Vol. 79, pp. 12-49, 1988.

[89] Osher, S., and Shu, C., *High-Order Nonoscillatory Schemes for Hamilton-Jacobi Equations*, Jour. Comp. Phys., Vol. 28, pp. 907-922, 1991.

[90] Perona, P., and Malik, J., *Scale-space and Edge Detection Using Anisotropic Diffusion*, IEEE Trans. Pattern Analysis and Machine Intelligence, Vol. 12(7), pp. 629–639, July 1990.

[91] Pindera, M.Z., and Talbot, L., *Flame-Induced Vorticity: The Effects of Stretch*, Twenty-First Symposium (Int'l) on Combustion, The Combustion Institute, Pittsburgh, pp. 1357-1366., 1986.

[92] Press, W. H., *Numerical Recipes*, Cambridge University Press, New York, 1988

[93] Puckett, E.G., *A Volume-of-Fluid Interface Tracking Algorithm with Applications to Computing Shock Wave Refraction*, Proceedings of the 4th International Symposium on Computational Computational Fluid Dynamics, Davis, California, 1991

[94] Rey, J.C.; Lie-Yea Cheng; McVittie, J.P.; Saraswat, K.C., *Monte Carlo Low Pressure Deposition Profile Simulations*, Journal of Vacuum Science and Technology A (Vacuum, Surfaces, and Films), May-June 1991, vol.9, (no.3, pt.1):1083-7.

[95] Rhee, C., Talbot, L., and Sethian, J.A., *Dynamical Study of a Premixed V flame*, Jour. Fluid Mech., 300, pp.87-115, 1995.

[96] Rouy, E. and Tourin, A., *A Viscosity Solutions Approach to Shape-From-Shading*, SIAM J. Num. Anal, 29, 3. pp. 867-884, 1992.

[97] Rudin, L., Osher, S., and Fatemi, E., *Nonlinear Total Variation-Based Noise Removal Algorithms*, Modelisations Matematiques pour le traitement d'images, IN-RIA, pp. 149–179, 1992.

[98] Sapiro, G., and Tannenbaum, A., *Area and Length Preserving Geometric Invariant Scale-spaces*, Proc. of Third European Conference on Computer Vision, LNCS Vol. 801, pp. 449–458, Stockholm, Sweden, May 1994.

[99] Sapiro, G., and Tannenbaum, A., *Affine Invariant Scale-Space*, Int. Jour. Comp. Vision, 11, 1, pp. 25-44, 1993.

[100] Sapiro, G., and Tannenbaum, A., *Image Smoothing Based on Affine Invariant Flow*, Proc. of the Conference on Information Sciences and Systems, Johns Hopkins University, March 1993.

[101] Scheckler, E.W., Ph.D. Dissertation, EECS, University of California, Berkeley, 1991.

[102] Scheckler, E.W., Toh, K.K.H., Hoffstetter, D.M., and Neureuther, A.R., *3D Lithography, Etching and Deposition Simulation*, Symposium on VLSI Technology, pp. 97-98, Oiso, Japan, (1991)

[103] Sedgewick, R., Algorithms, Addison-Wesley, 1988.

[104] Sethian, J.A., *Algorithms for Tracking Interfaces in CFD and Material Science*, Annual Review of Computational Fluid Mechanics, 1995

[105] Sethian, J.A., *An Analysis of Flame Propagation*, Ph.D. Dissertation, Mathematics, University of California, Berkeley, 1982.

[106] Sethian, J.A., *Curvature and the Evolution of Fronts*, Comm. in Mathematical Physics, Vol. 101, pp. 487–499, 1985.

[107] Sethian, J.A., *A Brief Overview of Vortex Methods*, in Vortex Methods and Vortex Motion, Eds. K. Gustafson and J.A. Sethian, SIAM Publications, Philadelphia, 1991.

[108] Sethian, J.A., *Curvature Flow and Entropy Conditions Applied to Grid Generation*, J. Comp. Phys. Vol. 115, pp.440-454, 1994.

[109] Sethian, J.A., *Level Set Techniques for Tracking Interfaces; Fast Algorithms, Multiple Regions, Grid Generation and Shape/Character Recognition*, in "Curvature Flows and Related Topics", Eds. A. Damlamian, J. Spruck, and A. Visintin, Gakuto Int. Series, Vol. 5., Tokyo, 1995.

[110] Sethian, J.A., *Level Set Methods: Evolving Interfaces in Geometry, Fluid Mechanics, Computer Vision and Material Science*, to appear, Cambridge University Press, 1996.

[111] Sethian, J.A., *A Marching Level Set Method for Monotonically Advancing Fronts*, Proc. Nat. Acad. Sci., Feb. 1996.

[112] Sethian, J.A., *Numerical Algorithms for Propagating Interfaces: Hamilton-Jacobi Equations and Conservation Laws*, Journal of Differential Geometry, Vol. 31, pp. 131–161, 1990.

[113] Sethian, J.A., *Numerical Methods for Propagating Fronts*, in Variational Methods for Free Surface Interfaces, edited by P. Concus and R. Finn, (Springer-Verlag, New Work, 1987).

[114] Sethian, J.A., *Parallel Level Set Methods for Propagating Interfaces on the Connection Machine*, Unpublished manuscript, 1989.

[115] Sethian, J.A., *Turbulent Combustion in Open and Closed Vessels*, J. Comp. Phys., Vol. 54, pp. 425-456 (1984).

[116] Sethian, J.A., Adalsteinsson, D., and Malladi, R., *Efficient Fast Marching Level Set Methods*, in progress, 1996.

[117] Sethian, J.A. and Strain, J.D., *Crystal Growth and Dendritic Solidification* J. Comp. Phys., Vol. 98, pp. 231-253, (1992).

[118] Souganidis, P.E., *Approximation Schemes for Viscosity Solutions of Hamilton-Jacobi Equations*, J. Diff. Eqns., 59, pp.1-43, 1985.

[119] Strain, J., *A Boundary Integral Approach to Unstable Solidification*, J. Comp. Phys., 85, pp. 342-389, (1989).

[120] Strain, J., *Linear Stability of Planar Solidification Fronts*, Physica D, 30, pp. 297-320, 1988

[121] Strain, J., *Velocity Effects in Unstable Solidification.* SIAM Jour. Appl. Math.,50,pp. 1-15. 1990.

[122] Sussman, M., Smereka, P. and Osher, S.J., *A Level Set Method for Computing Solutions to Incompressible Two-Phase Flow*, J. Comp. Phys. 114, pp.146-159, 1994.

[123] Taylor, J.E.; Cahn, J.W.; Handwerker, C.A., *Geometric models of crystal growth.*, Acta Metallurgica et Materialia, vol.40,7, pp. 1443-74, 1992.

[124] Terzopoulos, D., Witkin, A., and Kass, M., *Constraints on Deformable Models: Recovering 3D Shape and Nonrigid Motion*, Artificial Intelligence, 36, pp. 91–123, 1988.

[125] Toh, K.K.H., Ph.D. Dissertation, EECS, University of California, Berkeley, (1990).

[126] Toh, K.K.H., and Neureuther, A.R., *Three-Dimensional Simulation of Optical Lithography*, Proceedings SPIE, Optical/Laser Microlithography IV, vol. 1463, pp.356-367, (1991).

[127] Young, M.S., Lee, D., Lee., R., and Neureuther, A.R., *Extension of the Hopkins Theory of Partially Coherent Imaging to Include Thin-Film Interference effects* SPIE Optical/Laser Microlithography VI, vl 1927, pp. 452-463, (1993).

[128] Zhu, J., and Sethian, J.A., *Projection Methods Coupled to Level Set Interface Techniques*, J. Comp. Phys., vol. 102, pp.128-138, 1992.

[129] Zhu, J., and Ronney, P.D., *Simulation of Front Propagation at Large Non-dimensional Flow Disturbance Intensities*, to appear, Comb. Sci. Tech., 1995.